

08/30/00
JULY 13 U.S. PTO

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

10962 U.S. PTO
09/09/651425
08/30/00

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No. 003048.P008

Total Pages 3

First Named Inventor or Application Identifier Christopher Songer

Express Mail Label No. EL 234 217 656 US

ADDRESS TO: **Assistant Commissioner for Patents**
Box Patent Application
Washington, D. C. 20231

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. Fee Transmittal Form
(Submit an original, and a duplicate for fee processing)
2. Specification (Total Pages 39)
(preferred arrangement set forth below)
 - Descriptive Title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claims
 - Abstract of the Disclosure
3. Drawings(s) (35 USC 113) (Total Sheets 17)
4. Oath or Declaration (Total Pages 5)
 - a. Newly Executed (Original or Copy)
 - b. Copy from a Prior Application (37 CFR 1.63(d))
(for Continuation/Divisional with Box 17 completed) (Note Box 5 below)
 - i. **DELETIONS OF INVENTOR(S)** Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (useable if Box 4b is checked)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Microfiche Computer Program (Appendix)

7. Nucleotide and/or Amino Acid Sequence Submission
 (if applicable, all necessary)
 a. Computer Readable Copy
 b. Paper Copy (identical to computer copy)
 c. Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. Assignment Papers (cover sheet & documents(s))

9. a. 37 CFR 3.73(b) Statement (where there is an assignee)
 b. Power of Attorney (Executed)

10. English Translation Document (if applicable)

11. a. Information Disclosure Statement (IDS)/PTO-1449
 b. Copies of IDS Citations

12. Preliminary Amendment

13. Return Receipt Postcard (MPEP 503) (Should be specifically itemized)

14. a. Small Entity Statement(s)
 b. Statement filed in prior application, Status still proper and desired

15. Certified Copy of Priority Document(s) (if foreign priority is claimed)

16. Other: Attorney signature page including Copy of postcard and Certificate of Express Mailing pursuant to CFR §1.10.

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:

Continuation Divisional Continuation-in-part (CIP)
 of prior application No: _____

18. **Correspondence Address**

Customer Number or Bar Code Label (Insert Customer No. or Attach Bar Code Label here)
 or

Correspondence Address Below

NAME Dennis A. Nicholls, Reg. No. 42,036

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard

Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026

Country U.S.A. TELEPHONE (408) 720-8598 FAX (408) 720-9397

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: 30 August, 2000 By Dennis A. Nicholls
Dennis A. Nicholls
Reg. No. 42,036

12400 Wilshire Blvd.
Seventh Floor
Los Angeles, CA 90025-1026
(408) 720-8598

"Express Mail" mailing label number: EL 234 217 656 US

Date of Deposit: August 30, 2000

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231

Tina Domingo

(Typed or printed name of person mailing paper or fee)

Tina Domingo

(Signature of person mailing paper or fee)

8-30-2000

(Date signed)

Serial/Patent No.: * * Filing/Issue Date: Herewith
Client: Silicon Spice, Inc.
Title: System and Method For Preparing Software For Execution In A Dynamically Configurable Hardware Environment
BSTZ File No.: 03048.P008 Atty/Secty Initials: MES/DAN/td
Date Mailed: 8-30-2000 Docket Due Date: * *
The following has been received in the U.S. Patent & Trademark Office on the date stamped hereon:
 Amendment/Response (____ pgs.) Express Mail No.: EL234217656US Check No 37370
 Appeal Brief (____ pgs.) (in triplicate) _____ Month(s) Extension of Time Amt \$1,240.00
 Application - Utility (39 pgs., with cover and abstract) Information Disclosure Statement & PTO-144 (____ pgs.) Check No _____
 Application - Rule 1.53(b) Continuation (____ pgs.) Issue Fee Transmittal Amt. _____
 Application - Rule 1.53(b) Divisional (____ pgs.) Notice of Appeal _____
 Application - Rule 1.53(b) CIP (____ pgs.) Petition for Extension of Time _____
 Application - Rule 1.53(d) CPA Transmittal (____ pgs.) Petition for _____
 Application - Design (____ pgs.) Postcard _____
 Application - PCT (____ pgs.) Power of Attorney (____ pgs.)
 Application - Provisional (____ pgs.) Preliminary Amendment (____ pgs.)
 Assignment and Cover Sheet Reply Brief (____ pgs.)
 Certificate of Mailing (Express Mail) Response to Notice of Missing Parts
 Declaration & POA (5 pgs.) (Executed) Small Entity Declaration for Indep Inventor/Small Business
 Disclosure Docs & Org & Copy of Inventor Sgnd Letter (____ pgs.) Transmittal Letter, in duplicate (3 pages)
 Drawings: 17 # of sheets includes 17 figures Fee Transmittal, in duplicate
 Other _____

FEE CALCULATION (continued)**3. ADDITIONAL FEES**

<u>Large Entity</u>	<u>Small Entity</u>	<u>Fee Description</u>	<u>Fee Paid</u>
Fee Code	Fee (\$)	Fee Code (\$)	
105	130	205	65
127	50	227	25
			Surcharge - late filing fee or oath
			Surcharge - late provisional filing fee or cover sheet
139	130	139	130
147	2,520	147	2,520
112	920*	112	920*
			Non-English specification
			For filing a request for reexamination
			Requesting publication of SIR prior to Examiner action
113	1,840*	113	1,840*
			Requesting publication of SIR after Examiner action
115	110	215	55
116	380	216	190
117	870	217	435
118	1,360	218	680
128	1,850	228	925
119	300	219	150
120	300	220	150
121	260	221	130
138	1,510	138	1,510
140	110	240	55
			Notice of Appeal
			Filing a brief in support of an appeal
			Request for oral hearing
			Petition to institute a public use proceeding
			Petition to revive unavoidably abandoned application
141	1,210	241	605
			Petition to revive unintentionally abandoned application
142	1,210	242	605
143	430	243	215
144	580	244	290
122	130	122	130
123	50	123	50
126	240	126	240
581	40	581	40
			Petitions to the Commissioner
			Petitions related to provisional applications
			Submission of Information Disclosure Stmt
			Recording each patent assignment per property (times number of properties)
146	690	246	345
			For filing a submission after final rejection (see 37 CFR 1.129(a))
149	690	249	345
			For each additional invention to be examined (see 37 CFR 1.129(a))

Other fee (specify) _____

Other fee (specify) _____

SUBTOTAL (3) \$40.00

*Reduced by Basic Filing Fee Paid

SUBMITTED BY:

Typed or Printed Name: Dennis A. Nicholls
 Signature Dennis A. Nicholls Date 30 August, 2000
 Reg. Number 42,036 Deposit Account User ID _____
 (complete if applicable)

UNITED STATES PATENT APPLICATION

FOR

SYSTEM AND METHOD FOR
PREPARING SOFTWARE FOR EXECUTION IN A
DYNAMICALLY CONFIGURABLE HARDWARE
ENVIRONMENT

Inventors: Christopher Songer
Ian S. Eslick
Robert S. French

Prepared By:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN
12400 Wilshire Blvd., 7th Floor
Los Angeles, California 90025-1026
(408) 720-8300

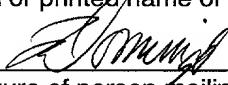
EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EL 234 217 656 US

Date of Deposit August 30, 2000

I hereby certify that I am causing this paper or fee to be
deposited with the United States Postal Service "Express Mail
Post Office to Addressee" service under 37 CFR 1.10 on the
date indicated above and is addressed to the Commissioner of
Patents and Trademarks, Washington, D.C. 20231

Tina Domingo
(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

8-30-2000
Date

SYSTEM AND METHOD FOR PREPARING SOFTWARE FOR EXECUTION
IN A DYNAMICALLY CONFIGURABLE HARDWARE ENVIRONMENT

FIELD OF THE INVENTION

5 The present invention relates to the field of software compilation and linking. In particular, the present invention relates to a system and method for preparing source code for execution in a dynamically configurable hardware environment.

BACKGROUND OF THE INVENTION

10 The software which executes upon processors is a sequence of digital words known as machine code. This machine code is understandable by the hardware of the processors. However, programmers typically write programs in a higher-level language which is much easier for humans to comprehend. The program listings in this higher-level language are called source code. In order 15 to convert the human-readable source code into machine-readable machine code, several special software tools are known in the art. These software tools are compilers, linkers, assemblers, and loaders.

Existing compilers, linkers, and assemblers prepare source code well in advance of their being executed upon processors. These software tools expect 20 that the hardware upon which the resulting machine code executes, including processors, will be in a predetermined and fixed configuration for the duration of the software execution. If a flexible processing methodology were invented, then the existing software tools would be inadequate to support processors and other hardware lacking a predetermined and fixed configuration.

SUMMARY OF THE INVENTION

A system and method for creating run time executables for a configurable processing element array is disclosed. This system and method includes the step of partitioning a processing element array into a number of hardware accelerators, which in one embodiment are called bins. The system and method then involves decomposing a program description into a number of kernel sections. Next, mapping the kernel sections into a number of hardware dependent designs is performed. Finally, a matrix of the hardware accelerators, which may include bins, and the designs is formed for use by the run time system.

003048.P008

BRIEF DESCRIPTION OF THE DRAWINGS

The features, aspects, and advantages of the present invention will become more fully apparent from the following detailed description, appended claims, and accompanying drawings in which:

5 Figure 1 is the overall chip architecture of one embodiment. This chip architecture comprises many highly integrated components.

Figure 2 is an eight bit multiple context processing element (MCPE) core of one embodiment of the present invention.

Figure 3 is a data flow diagram of the MCPE of one embodiment.

10 Figure 4 shows the major components of the MCPE control logic structure of one embodiment.

Figure 5 is the finite state machine (FSM) of the MCPE configuration controller of one embodiment.

15 Figure 6 is a data flow system diagram of the preparation of run time systems tables by the temporal automatic place and route (TAPR) of one embodiment.

Figure 7A is a block diagram of a system including exemplary MCPEs, according to one embodiment.

20 Figure 7B is a block diagram of a system including exemplary digital signal processors (DSP), according to one embodiment.

Figure 8 is a diagram of the contents of an exemplary run time kernel (RTK), according to one embodiment.

Figure 9A is a process chart showing the mapping of an exemplary single threaded process into kernel segments, according to one embodiment.

Figure 9B is a process chart showing the allocation of the kernel segments of Figure 9A into multiple bins.

5 Figure 9C is a process chart showing the allocation of the kernel segments of two processes into multiple bins.

Figure 10 is an exemplary TAPR table, according to one embodiment.

Figure 11 is a diagram of a first exemplary variant of a design, according to one embodiment.

10 Figure 12 is a diagram of a second exemplary variant of a design, according to another embodiment.

Figure 13 is a diagram of an exemplary logical MCPE architecture, according to one embodiment.

15 Figure 14 is a diagram of an exemplary logical processor-based architecture, according to one embodiment.

DETAILED DESCRIPTION OF THE INVENTION

In the following description, numerous specific details are set forth to provide a thorough understanding of the present invention. However, one having an ordinary skill in the art may be able to practice the invention without 5 these specific details. In some instances, well-known circuits, structures, and techniques have not been shown in detail to not unnecessarily obscure the present invention.

Figure 1 is the overall chip architecture of one embodiment. This chip architecture comprises many highly integrated components. While prior art 10 chip architectures fix resources at fabrication time, specifically instruction source and distribution, the chip architecture of the present invention is flexible. This architecture uses flexible instruction distribution that allows position independent configuration and control of a number of multiple context processing elements (MCPEs) resulting in superior performance provided by the 15 MCPEs. The flexible architecture of the present invention uses local and global control to provide selective configuration and control of each MCPE in an array; the selective configuration and control occurs concurrently with present function execution in the MCPEs.

The chip of one embodiment of the present invention is composed of, but 20 not limited to, a 10x10 array of identical eight-bit functional units, or MCPEs 102, which are connected through a reconfigurable interconnect network. The MCPEs 102 serve as building blocks out of which a wide variety of computing structures may be created. The array size may vary between 2x2 MCPEs and

16x16 MCPEs, or even more depending upon the allowable die area and the desired performance. A perimeter network ring, or a ring of network wires and switches that surrounds the core array, provides the interconnections between the MCPEs and perimeter functional blocks.

5 Surrounding the array are several specialized units that may perform functions that are too difficult or expensive to decompose into the array. These specialized units may be coupled to the array using selected MCPEs from the array. These specialized units can include large memory blocks called configurable memory blocks 104. In one embodiment these configurable
10 memory blocks 104 comprise eight blocks, two per side, of 4 kilobyte memory blocks. Other specialized units include at least one configurable instruction decoder 106.

15 Furthermore, the perimeter area holds the various interfaces that the chip of one embodiment uses to communicate with the outside world including:
20 input/output (I/O) ports; a peripheral component interface (PCI) controller, which may be a standard 32-bit PCI interface; one or more synchronous burst static random access memory (SRAM) controllers; a programming controller that is the boot-up and master control block for the configuration network; a master clock input and phase-locked loop (PLL) control/configuration; a Joint Test Action Group (JTAG) test access port connected to all the serial scan chains on the chip; and I/O pins that are the actual pins that connect to the outside world.

Two concepts which will be used to a great extent in the following description are context and configuration. Generally, “context” refers to the definition of what hardware registers in the hardware perform which function at a given point in time. In different contexts, the hardware may perform

5 differently. A bit or bits in the registers may define which definition is currently active. Similarly, “configuration” usually refers to the software bits that command the hardware to enter into a particular context. This set of software bits may reside in a register and define the hardware’s behavior when a particular context is set.

10 Figure 2 is an eight bit MCPE core of one embodiment of the present invention. Primarily the MCPE core comprises memory block 210 and basic ALU core 220. The main memory block 210 is a 256 word by eight bit wide memory, which is arranged to be used in either single or dual port modes. In dual port mode the memory size is reduced to 128 words in order to be able to 15 perform two simultaneous read operations without increasing the read latency of the memory. Network port A 222, network port B 224, ALU function port 232, control logic 214 and 234, and memory function port 212 each have configuration memories (not shown) associated with them. The configuration memories of these elements are distributed and are coupled to a Configuration

20 Network Interface (CNI) (not shown) in one embodiment. These connections may be serial connections but are not so limited. The CNI couples all configuration memories associated with network port A 222, network port B 224, ALU function port 232, control logic 214 and 234, and memory function

port 212 thereby controlling these configuration memories. The distributed configuration memory stores configuration words that control the configuration of the interconnections. The configuration memory also stores configuration information for the control architecture. Optionally it can also be a multiple
5 context memory that receives context selecting signals which have been broadcast globally and locally from a variety of sources.

Figure 3 is a data flow diagram of the MCPE of one embodiment. The structure of each MCPE allows for a great deal of flexibility when using the MCPEs to create networked processing structures. The major components of

10 the MCPE include static random access memory (SRAM) main memory 302, ALU with multiplier and accumulate unit 304, network ports 306, and control logic 308. The solid lines mark data flow paths while the dashed lines mark control paths; all of the lines are one or more bits wide in one embodiment. There is a great deal of flexibility available within the MCPE because most of
15 the major components may serve several different functions depending on the MCPE configuration.

The MCPE main memory 302 is a group of 256 eight bit SRAM cells that can operate in one of four modes. It takes in up to two eight bit addresses from A and B address/data ports, depending upon the mode of operation. It also
20 takes in up to four bytes of data, which can be from four floating ports, the B address/data port, the ALU output, or the high byte from the multiplier. The main memory 302 outputs up to four bytes of data. Two of these bytes, memory A and B, are available to the MCPE's ALU and can also be directly

driven onto the level 2 network. The other two bytes, memory C and D, are only available to the network. The output of the memory function port 306 controls the cycle-by-cycle operation of the memory 302 and the internal MCPE data paths as well as the operation of some parts of the ALU 304 and the 5 control logic 308. The MCPE main memory may also be implemented as a static register file in order to save power.

Each MCPE contains a computational unit 304 comprised of three semi-independent functional blocks. The three semi-independent functional blocks comprise an eight bit wide ALU, an 8x8 to sixteen bit multiplier, and a sixteen 10 bit accumulator. The ALU block, in one embodiment, performs logical, shift, arithmetic, and multiplication operations, but is not so limited. The ALU function port 306 specifies the cycle-by-cycle operation of the computational unit. The computational units in orthogonally adjacent MCPEs can be chained 15 to form wider-word data paths.

The MCPE network ports 306 connect the MCPE network to the internal MCPE logic (memory, ALU, and control). There are eight network ports 306 in each MCPE, each serving a different set of purposes. The eight network ports 306 comprise two address/data ports, two function ports, and four floating 20 ports. The two address/data ports feed addresses and data into the MCPE memories and ALU. The two function ports feed instructions into the MCPE logic. The four floating ports may serve multiple functions. The determination of what function they are serving is made by the configuration of the receivers of their data.

The MCPEs of one embodiment are the building blocks out of which more complex processing structures may be created. The structure that joins the MCPE cores into a complete array in one embodiment is actually a set of several mesh-like interconnect structures. Each interconnect structure forms

5 a network, and each network is independent in that it uses different paths, but the networks do join at the MCPE input switches. The network structure of one embodiment of the present invention is comprised of a local area broadcast network (level 1), a switched interconnect network (level 2), a shared bus network (level 3), and a broadcast, or configuration, network.

10 Figure 4 shows the major components of the MCPE control logic structure of one embodiment. The Control Tester 602 takes the output of the ALU or two bytes from floating ports 604 and 606, plus the left and right carryout bits, and performs a configurable test on them. The result is one bit indicating if the comparison matched. This bit is referred to as the control bit.

15 This Control Tester 602 serves two main purposes. First, it acts as a programmable condition code generator testing the ALU output for any condition that the application needs to test for. Secondly, since these control bits can be grouped and sent out across the level 2 and 3 networks, this unit can be used to perform a second or later stage reduction on a set of control

20 bits/data generated by other MCPE's.

The level 1 network 608 carries the control bits. The level 1 network 608 consists of direct point-to-point communications between every MCPE and its 12 nearest neighbors. Thus, each MCPE will receive 13 control bits (12

neighbors and its own) from the level 1 network. These 13 control bits are fed into the Control Reduce block 610 and the BFU input ports 612. The Control Reduce block 610 allows the control information to rapidly effect neighboring MCPEs. The MCPE input ports allow the application to send the control data

5 across the normal network wires so they can cover long distances. In addition the control bits can be fed into MCPEs so they can be manipulated as normal data.

The Control Reduce block 610 performs a simple selection on either the control words coming from the level 1 control network, the level 3 network, or

10 two of the floating ports. The selection control is part of the MCPE configuration. The Control Reduce block 610 selection results in the output of five bits. Two of the output bits are fed into the MCPE configuration controller 614. One output bit is made available to the level 1 network, and one output bit is made available to the level 3 network.

15 The MCPE configuration controller 614 selects on a cycle-by-cycle basis which context, major or minor, will control the MCPE's activities. The

controller consists of a finite state machine (FSM) that is an active controller and not just a lookup table. The FSM allows a combination of local and global control over time that changes. This means that an application may run for a

20 period based on the local control of the FSM while receiving global control signals that reconfigure the MCPE, or a block of MCPEs, to perform different functions during the next clock cycle. The FSM provides for local configuration and control by locally maintaining a current configuration context for control of

the MCPE. The FSM provides for global configuration and control by providing the ability to multiplex and change between different configuration contexts of the MCPE on each different clock cycle in response to signals broadcast over a network. This configuration and control of the MCPE is powerful because it

5 allows an MCPE to maintain control during each clock cycle based on a locally maintained configuration context while providing for concurrent global on-the-fly reconfiguration of each MCPE. This architecture significantly changes the area impact and characterization of an MCPE array while increasing the efficiency of the array without wasting other MCPEs to perform the

10 configuration and control functions.

Figure 5 is the FSM 502 of the MCPE configuration controller of one embodiment. In controlling the functioning of the MCPE, control information 504 is received by the FSM 502 in the form of state information from at least one surrounding MCPE in the networked array. This control information is in

15 the form of two bits received from the Control Reduce block of the MCPE control logic structure. In one embodiment, the FSM 502 also has three state bits that directly control the major and minor configuration contexts for the particular MCPE. The FSM 502 maintains the data of the current MCPE configuration by using a feedback path 506 to feed back the current

20 configuration state of the MCPE of the most recent clock cycle. The feedback path 506 is not limited to a single path. The FSM 502 selects one of the available configuration memory contexts for use by the corresponding MCPE during the next clock cycle in response to the received state information from

the surrounding MCPEs and the current configuration data. This selection is output from the FSM 502 in the form of a configuration control signal 508.

The selection of a configuration memory context for use during the next clock cycle occurs, in one embodiment, during the execution of the configuration

5 memory context selected for the current clock cycle.

Figure 6 is a data flow system diagram of the preparation of run time systems tables by the temporal automatic place and route (TAPR) of one embodiment. In step 650 an application program in source code is selected. In the Figure 6 embodiment the application program is written in a procedural oriented language, C, but in other embodiments the application program could be written in another procedural oriented language, in an object oriented language, or in a dataflow language.

The source code of step 650 is examined in decision step 652. Portions of the source code are separated into overhead code and kernel code sections.

15 Kernel code sections are defined as those routines in the source code which may be advantageously executed in a hardware accelerator. Overhead code is defined as the remainder of the source code after all the kernel code sections are identified and removed.

In one embodiment, the separation of step 652 is performed by a software profiler. The software profiler breaks the source code into functions. In one embodiment, the complete source code is compiled and then executed with a representative set of test data. The profiler monitors the timing of the execution, and then based upon this monitoring determines the function or

functions whose execution consumes a significant portion of execution time.

Profiler data from this test run may be sent to the decision step 652. The profiler identifies these functions as kernel code sections.

In an alternate embodiment, the profiler examines the code of the

- 5 functions and then identifies a small number of functions that are anticipated to consume a large portion of the execution runtime of the source code. These functions may be identified by attributes such as having a regular structure, having intensive mathematical operations, having a repeated or looped structure, and having a limited number of inputs and outputs. Attributes
- 10 which argue against the function being identified as kernel sections include numerous branches and overly complex control code.

In an alternate embodiment, the compiler examines the code of the functions to determine the size of arrays traversed and the number of variables that are live during the execution of a particular block or function. Code that

- 15 has less total memory used than that in the hardware accelerators and associated memories are classified as kernel code sections. The compiler may use well-understood optimization methods such as constant propagation, loop induction, in-lining and intra-procedural value range analysis to infer this information from the source code.

- 20 Those functions that are identified as kernel code section by one of the above embodiments of profiler, are then labeled, in step 654, as kernel code sections. The remainder of the source code is labeled as overhead code. In

alternate embodiments, the separation of step 652 may be performed manually by a programmer.

In step 656, the Figure 6 process creates hardware designs for implementing the kernel code sections of step 654. These designs are the executable code derived from the source code of the kernel code sections. Additionally, the designs contain any necessary microcode or other fixed-constant values required in order to run the executable code on the target hardware. The designs are not compiled in the traditional sense. Instead they are created by the process of step 656 which allows for several embodiments.

10 In one embodiment, the source code of the kernel code section is compiled automatically by one of several compilers corresponding to the available hardware accelerators. In an alternate embodiment, a programmer may manually realize the executable code from the source code of the kernel code sections, as shown by the dashed line from step 656 to step 550. In a 15 third embodiment the source code of the kernel code sections is compiled automatically for execution on both the processors and the hardware accelerators, and both versions are loaded into the resulting binary. In a fourth embodiment, a hardware accelerator is synthesized into a custom hardware accelerator description.

20 In step 658 the hardware designs of step 656 are mapped to all available target hardware accelerators. The target hardware accelerators may be a processor (such as a digital signal processor or DSP), an MCPE, or a defined set of MCPEs called a bin. A bin may contain any number of MCPEs from one to

the maximum number of MCPEs on a given integrated circuit. However, in one embodiment a quantity of 12 MCPEs per bin is used. The MCPEs in each bin may be geometrically neighboring MCPEs, or the MCPEs may be distributed across the integrated circuit. However, in one embodiment the MCPEs of each 5 bin are geometrically neighboring.

In the temporal automatic place and route (TAPR) of step 660, the microcode created in step 656 may be segmented into differing context-dependent portions. For example, a given microcode design may be capable of loading and executing in either lower memory or upper memory of a given bin.

10 The TAPR of step 660 may perform the segmentation in several different ways depending upon the microcode. If, for example, the microcode is flat, then the microcode may only be loaded into memory in one manner. Here no segmentation is possible. Without segmentation one microcode may not be background loaded onto a bin's memory. The bin must be stalled and the 15 microcode loaded off-line.

In another example, memory is a resource which may be controlled by the configuration. It is possible for the TAPR of step 660 to segment microcode into portions, corresponding to differing variants, which correspond to differing contexts. For example, call one segmented microcode portion context 2 and 20 another one context 3. Due to the software separation of the memory of the bin it would be possible to place the context 2 and context 3 portions into lower memory and upper memory, respectively. This allows background loading of one portion while another portion is executing.

The TAPR of step 660 supports two subsequent steps in the preparation of the source code for execution. In step 664, a table is prepared for subsequent use by the run time system. In one embodiment, the table of step 664 contains all of the three-tuples corresponding to allowable combinations of 5 designs (from step 656), bins, and variants. A variant of a design or a bin is any differing implementation where the functional inputs and the outputs are identical when viewed from outside. The variants of step 664 may be variants of memory separation, such as the separation of memory into upper and lower memory as discussed above. Other variants may include differing geometric 10 layouts of MCPEs within a bin, causing differing amounts of clock delays being introduced into the microcodes, and also whether or not the use of various parts of the MCPEs overlap the use of the parts in other variants. In each case a variant performs a function whose inputs and outputs are identical outside of the function. The entries in the table of step 664 point to executable binaries, 15 each of which may each be taken and executed without further processing at run time. The table of step 664 is a set of all alternative execution methods available to the run time system for a given kernel section.

The other step supported by the TAPR of step 660 is the creation of configurations, microcodes, and constants of step 662. These are the 20 executable binaries which are pointed to by the entries in the table of step 664.

Returning now to decision step 652, the portions of the source code which were previously deemed overhead are sent to a traditional compiler 670 for compilation of object code to be executed on a traditional processor.

Alternately, the user may hand code the source program into the assembly language of the target processor. The overhead C code may also be nothing more than calls to kernel sections. The object code is used to create object code files at step 672.

5 Finally, the object code files of step 672, the configurations, microcode, and constants of step 662, and table of step 664 are placed together in a format usable by the run time system by the system linker of step 674.

Note that the instructions for the process of Figure 6 may be described in software contained in a machine-readable medium. A machine-readable

10 medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g. a computer). For example, a machine-readable medium includes read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; and electrical, optical, acoustical, or other form of propagated signals (e.g. carrier waves, infrared signals, digital signals, etc.).

15 Figure 7A is a block diagram of a system including exemplary MCPEs, according to one embodiment. Chip architecture 700 includes processing elements processor A 702, processor B 720, bin 0 706, bin 1 708, and bin 2 710. In the Figure 7A embodiment, the function of hardware accelerator may 20 be assigned to the MCPEs, either individually or grouped into bins. A run-time kernel (RTK) 704 apportions the executable software among these processing elements at the time of execution. In the Figure 7A embodiment, processor A 702 or processor B 720 may execute the overhead code identified in step 652

and created as object files in step 672 of the Figure 6 process. Bin 0 706, bin 1 708, and bin 2 710 may execute the kernel code identified in step 652.

Each processing element processor A 702 and processor B 720 is supplied with an instruction port, instruction port 724 and instruction port 5 722, respectively, for fetching instructions for execution of overhead code.

Bin 0 706, bin 1 708, and bin 2 710 contain several MCPEs. In one embodiment, each bin contains 12 MCPEs. In alternate embodiments, the bins could contain other numbers of MCPEs, and each bin could contain a different number of MCPEs than the other bins.

10 In the Figure 7A embodiment, bin 0 706, bin 1 708, and bin 2 710 do not share any MCPEs, and are therefore called non-overlapping bins. In other embodiments, bins may share MCPEs. Bins which share MCPEs are called overlapping bins.

RTK 704 is a specialized microprocessor for controlling the configuration 15 of chip architecture 700 and controlling the loading and execution of software in bin 0 706, bin 1 708, and bin 2 710. In one embodiment, RTK 704 may move data from data storage 728 and configuration microcode from configuration microcode storage 726 into bin 0 706, bin 1 708, and bin 2 710 in accordance with the table 730 stored in a portion of data storage 728. In 20 alternate embodiments, RTK 704 may move data from data storage 728, without moving any configuration microcode from configuration microcode storage 726. Table 730 is comparable to that table created in step 664 discussed in connection with Figure 6 above.

Figure 7B is a block diagram of a system including exemplary digital signal processors (DSP), according to one embodiment. Chip architecture 750 includes processing elements processor A 752, processor B 770, DSP 0 756, DSP 1 758, and DSP 2 760. In the Figure 7B embodiment, the function of 5 hardware accelerator may be assigned to the DSPs. In other embodiments, DSP 0 756, DSP 1 758, and DSP 2 760 may be replaced by other forms of processing cores. A run-time kernel (RTK) 754 apportions the executable software among these processing elements at the time of execution.

In the Figure 7B embodiment, processor A 752 or processor B 770 may 10 execute the overhead code identified in step 652 and created as object files in step 672 of the Figure 6 process. DSP 0 756, DSP 1 758, and DSP 2 760 may execute the kernel code identified in step 652. Each processing element processor A 702 and processor B 720 is supplied with an instruction port, instruction port 724 and instruction port 722, respectively, for fetching 15 instructions for execution of overhead code.

One difference between the Figure 7A and Figure 7B embodiments is that the Figure 7B embodiment lacks an equivalent to the configuration microcode storage 726 of Figure 7A. No configuration microcode is required as the DSPs of Figure 7B have a fixed instruction set (microcode) architecture.

20 RTK 754 is a specialized microprocessor for controlling the configuration of chip architecture 750 and controlling the loading and execution of software in DSP 0 756, DSP 1 758, and DSP 2 760. In one embodiment, RTK 754 may move data from data storage 778 into DSP 0 756, DSP 1 758, and DSP 2 760 in

accordance with the table 780 stored in a portion of data storage 778. Table 780 is comparable to that table created in step 664 discussed in connection with Figure 6 above.

Figure 8 is a diagram of the contents of an exemplary run time kernel 5 (RTK) 704, according to one embodiment. RTK 704 contains several functions in microcontroller form. In one embodiment, these functions include configuration direct memory access (DMA) 802, microcode DMA 804, arguments DMA 806, results DMA 808, and configuration network source 810. RTK 704 utilizes these functions to manage the loading and execution of kernel 10 code and overhead code on chip architecture 700. Configuration DMA 802, microcode DMA 804, arguments DMA 806, and results DMA 808 each comprise a simple hardware engine for reading from one memory and writing to another.

Configuration DMA 802 writes configuration data created by the TAPR 15 660 in step 622 of the Figure 6 process. This configuration data configures a bin to implement the behavior of the kernel code section determined in the table-making step 664 of Figure 6. The configuration data transfers are under the control of RTK 704 and the configuration data itself is entered in table 730. Configuration data is unchanged over the execution of the hardware 20 accelerator.

Microcode DMA 804 writes microcode data for each configuration into the bins. This microcode further configures the MCPEs with instruction data that allows the function of the hardware accelerator to be changed on a cycle-

by-cycle basis while the hardware accelerator is executing. Each bin may have multiple microcode data sets available for use. Microcode data is stored in the configuration microcode storage 726 and written into memory within the MCPEs of each bin by microcode DMA 804.

5 Arguments DMA 806 and results DMA 808 set up transfers of data from data memory 728 into one of the bins bin 0 706, bin 1 708, or bin 2 710. Argument data are data stored in a memory by a general purpose processor which requires subsequent processing in a hardware accelerator. The argument data may be considered the input data of the kernel code sections 10 executed by the bins. Results data are data sent from the hardware accelerator to the general purpose processor as the end product of a particular kernel code section's execution in a bin. The functional units arguments DMA 806 and results DMA 808 transfer this data without additional processor intervention.

Configuration network source 810 controls the configuration network. 15 The configuration network effects the configuration of the MCPEs of the bins bin 0 70, bin 1 708 and bin 2 710, and of the level 1, level 2, and level 3 interconnect described in Figure 3 and Figure 4. Configuration of the networks enables the RTK to control the transfer of configuration data, microcode data, arguments data, and results data amongst the data memory 728, configuration 20 memory 726, and the MCPEs of bin 0 706, bin 1 708, and bin 2 710.

In cases where there are multiple contexts, RTK 704 may perform background loading of microcode and other data while the bins are executing

kernel code. An example of this is discussed below in connection with Figure 11.

Figure 9A is a process chart showing the mapping of an exemplary single threaded process into kernel segments, according to one embodiment.

5 Source code 1 900 and source code 2 960 are two exemplary single threaded processes which may be used as the C source code 650 of the Figure 6 process. In one embodiment, source code 1 900 may contain overhead code 910, 914, 918, 922, 926, and 930, as well as kernel code 912, 916, 920, 924, and 928. The identification of the overhead code and kernel code sections may be 10 performed in step 652 of the Figure 6 process. Overhead code 910, 914, 918, 922, 926, and 930 may be executed in processor A 702 or processor B 720 of the Figure 7A embodiment. Kernel code 912, 916, 920, 924, and 928 may be executed in bin 0 706, bin 1 708, or bin 2 710 of the Figure 7A embodiment. The TAPR 660 of the Figure 6 process may create the necessary configurations 15 and microcode for the execution of the kernel code 912, 916, 920, 924, and 928.

Figure 9B is a process chart showing the allocation of the kernel segments of Figure 9A into multiple bins. Utilizing the table 780 produced in step 664 of the Figure 6 process, RTK 704 may load and execute the overhead 20 code 910, 914, 918, 922, 926, and 930 and the kernel code 912, 916, 920, 924, and 928 into an available processor or bin as needed. In the exemplary Figure 9B embodiment, RTK 704 loads the first overhead code 910 into

processor A 702 for execution during time period 970. RTK 704 then loads the first kernel code 912 into bin 0 706 for execution during time period 972.

Depending upon whether overhead code 914 requires the completion of kernel code 912, RTK 704 may load overhead code 914 into processor A 702 for 5 execution during time period 974. Similarly, depending upon whether kernel code 916 requires the completion of overhead code 914 or kernel code 910, RTK 704 may load kernel code 916 into bin 1 708 for execution during time period 976.

Depending upon requirements for completion, RTK 704 may continue to 10 load and execute the overhead code and kernel code in an overlapping manner in the processors and the bins. When overhead code or kernel code require the completion of a previous overhead code or kernel code, RTK 704 may load the subsequent overhead code or kernel code but delay execution until the required completion.

15 Figure 9C is a process chart showing the allocation of the kernel segments of two processes into multiple bins. In the Figure 9C embodiment, source code 1 900 and source code 2 960 may be the two exemplary single threaded processes of Figure 9A. Prior to the execution of source code 1 900 and source code 2 960 in Figure 9C, the kernel code and overhead code 20 sections may be identified and processed in the Figure 6 process or in an equivalent alternate embodiment process. Utilizing the table 730 for source code 1 900, produced in step 664 of the Figure 6 process, RTK 704 may load and execute the overhead code 910, 914, 918, and 922, and the kernel code

912, 916, and 920 into an available processor or bin as needed. Similarly, an equivalent table (not shown) may be prepared for source code 2 960. In the Figure 9C embodiment, by utilizing this equivalent table for source code 2 960, RTK 704 may load and execute the overhead code 950, 954, and 958, and the 5 kernel code 952 and 956, into an available processor or bin as needed.

In the exemplary Figure 9C embodiment, RTK 704 loads the first overhead code 910, 960 sections into processor A 702 and processor B 720, respectively, for execution in time periods 980 and 962, respectively.

When overhead code 910 finishes executing, RTK 704 may load kernel 10 code 912 into bin 0 706 for execution in time period 982. When kernel code 912 finishes executing, RTK 704 may load the next overhead code 914 into an available processor such as processor B 720 during time period 948.

When overhead code 950 finishes executing, RTK 704 may load kernel 15 code 952 into available bin 1 708 for execution during time period 964. When kernel code 952 finishes executing RTK 704 may load the next overhead code 954 into processor A 702 for execution during time period 966.

Therefore, as shown in Figure 9C, multiple threads may be executed utilizing the designs, bins, and tables of various embodiments of the present invention. The overhead code and kernel code sections of the several threads 20 may be loaded and executed in an overlapping manner among the several processors and bins available.

Figure 10 is an exemplary TAPR table, according to one embodiment. The TAPR table of Figure 10 is a three dimensional table, containing entries

that are three-tuples of the possible combinations of bins, designs, and variants. The TAPR table contains more than just a recitation of the designs of the kernel code segments mapped into the bins (hardware accelerators).

Instead, the TAPR table includes the dimension of variants of the bins. Each

5 combination of designs and bins may have multiple variants. Variants perform the identical function from the viewpoint of the inputs and outputs, but differ in implementation. An example is when bins are configured from a 3 by 4 array of MCPEs as versus a 4 by 3 array of MCPEs. In this case differing timing requirements due to differing path lengths may require separate

10 variants in the configuration and microcode data of the hardware accelerator.

In one embodiment, these variants may take the form of different microcode implementations of the design, or the variants may be differing signal routing paths among the MCPEs of the bins. Two additional exemplary variants are discussed below in connection with Figure 11 and Figure 12.

15 In the Figure 10 embodiment, it should be noted that it is not necessary that there be an valid entry at each location of the matrix. There are situations where there are relatively few valid sets of designs, bins, and variants. These situations give rise to a sparsely-populated TAPR matrix. In other situations, there may be a valid set of designs, bins and variants for all locations in the matrix. These situations give rise to a fully-populated TAPR matrix.

Figure 11 is a diagram of a first exemplary variant of a design, according to one embodiment. Memory available to a bin is a resource that may be controlled by the configuration. In this embodiment, bin 0 706 may have a

memory that is logically partitioned into a lower memory 1104 and an upper memory 1102. Each memory area, for example upper memory 1102 and lower memory 1104, may be running a different context. For example, there could be a context 2 running in upper memory 1102 and an alternate context 3 loaded

5 in lower memory 1104.

Bin 0 706 is configured in accordance with a design, but depending upon how the design is loaded in memory certain instructions such as jump and load may have absolute addresses embedded in them. Therefore the design may have a variant for loading in upper memory 1102 under the control of

10 context 2 and a second variant for loading in lower memory 1104 under the control of context 3. Having multiple variants in this manner advantageously allows any run-time engine such as RTK 704 to load the microcode for one variant in either upper memory 1102 or lower memory 1104 while execution is still proceeding in the alternate memory space under a different context.

15 Figure 12 is a diagram of a second exemplary variant of a design, according to another embodiment. The memory available to bin 1 708 may be in two physically distinct areas on the chip. In Figure 12 one section of memory may be at physical location 1202 with data path 1212, and another section of memory may be at physical location 1204 with data path 1214. If

20 data path 1214 is physically longer than data path 1212 then it may be necessary to insert additional clock cycles for a given design to run on bin 1 708 from memory at physical location 1202 in comparison with physical

location 1204. Here the two variants differ in the number of internal wait states in the microcode of the design.

Figure 13 is a diagram of an exemplary logical MCPE architecture 1300, according to one embodiment. Included within architecture 1300 are main processor 1304, run time kernel (RTK) processor 1316, an instruction memory (IMEM) 1302, a processor data memory 1306 with attached DMA 1308, and a configuration memory 1310 with attached DMA 1312. RTK processor 1316 is connected to a control bus 1314, which controls the operation of DMA 1308 and DMA 1312. DMA 1308 in turn generates an argument bus 1318, and DMA 1312 in turn generates a configuration bus 1328.

Architecture 1300 also includes several hardware accelerators 1320, 1330, 1340. Each accelerator contains a local DMA for sending and receiving data to and from the argument bus 1318 and a DMA for receiving data from the configuration bus 1328. For example, accelerator 1320 has DMA 1322 for sending and receiving data to and from the argument bus 1318 and DMA 1324 for receiving data from the configuration bus 1328. In the Figure 13 embodiment, argument bus 1318 is a bi-directional bus that may carry instruction data, argument data, and results data.

Figure 14 is a diagram of an exemplary logical processor-based architecture, according to one embodiment. Included within architecture 1400 are main processor 1404, run time kernel (RTK) processor 1416, an instruction memory (IMEM) 1402 with attached DMA 1412, and a processor data memory 1406 with attached DMA 1408. RTK processor 1416 generates a control bus

1414, which controls the operation of DMA 1408, 1412. DMA 1408 in turn generates an argument bus 1418, and DMA 1412 in turn generates an instruction bus 1428.

Architecture 1400 also includes several DSPs 1420, 1430, 1440. Each

- 5 DSP is connected to a DMA controller for receiving argument data from the argument bus 1418 and a data cache for temporary storage of the argument data. Each DSP is also connected to a DMA controller for receiving instruction data from the instruction bus 1418 and an instruction cache for temporary storage of the instruction data. Both sets of DMA controller receive control
- 10 from the control bus 1414. For example, DSP 1420 has DMA controller 1428 for receiving data from the argument bus 1418 and data cache 1426 for temporary storage of the argument data. DSP 1420 also has DMA controller 1422 for receiving data from the instruction bus 1428 and instruction cache 1424 for temporary storage of the instruction data. In the Figure 14
- 15 embodiment, argument bus 1418 carries argument data but does not carry instruction data.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will however be evident that various modifications and changes can be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. Therefore, the scope of the invention should be limited only by the appended claims.

CLAIMS

What is claimed is:

- 1 1. A method of creating run time executable code, comprising:
 - 2 partitioning a processing element array into a plurality of hardware
 - 3 accelerators;
 - 4 decomposing a program description into a plurality of kernel sections;
 - 5 mapping said kernel sections into a plurality of hardware dependent
 - 6 designs; and
 - 7 forming a matrix describing said hardware accelerators and said designs
 - 8 configured to support run time execution.
- 1 2. The method of claim 1, wherein said partitioning includes partitioning
- 2 into digital signal processors.
- 1 3. The method of claim 1, wherein said partitioning includes partitioning
- 2 into bins.
- 1 4. The method of claim 1, wherein said mapping includes mapping into
- 2 multiple hardware contexts.

1 5. The method of claim 4, wherein said mapping into multiple hardware
2 contexts includes mapping a first set of variants.

1 6. The method of claim 5, wherein said first set of variants are produced
2 based upon resource usage.

1 7. The method of claim 5, wherein said mapping includes mapping a second
2 set of variants of said designs configured to support multiple hardware
3 configurations of one of a plurality of bins.

1 8. The method of claim 1, wherein said mapping is performed by a place
2 and route.

1 9. The method of claim 1, wherein said decomposing is performed
2 manually.

1 10. The method of claim 1, wherein said decomposing is performed by a
2 software profiler.

1 11. The method of claim 10, wherein said decomposing includes executing
2 code compiled from said program description and monitoring timing of said
3 executing.

1 12. The method of claim 11, wherein said executing utilizes a set of test data.

1 13. The method of claim 11, wherein said monitoring includes determining
2 functions that consume a significant portion of said timing of said executing.

1 14. The method of claim 10, wherein said decomposing includes identifying
2 kernel sections by identifying regular structures.

1 15. The method of claim 10, wherein said decomposing includes identifying
2 kernel sections by identifying sections with a limited number of inputs and
3 outputs.

1 16. The method of claim 10, wherein said decomposing includes identifying
2 kernel sections by identifying sections with a limited number of branches.

1 17. The method of claim 10, wherein decomposing identifies overhead
2 sections.

1 18. The method of claim 1, wherein mapping includes creating microcode.

1 19. The method of claim 1, wherein said mapping includes creating context
2 dependent configurations.

1 20. The method of claim 1, wherein said matrix is sparsely-populated.

1 21. The method of claim 1, wherein said matrix is fully-populated.

1 22. A system for creating run time executable code, comprising:

2 a plurality of hardware accelerators partitioned from a processing
3 element array;

4 a plurality of kernel sections created from a program description;

5 a plurality of hardware dependent designs derived from said kernel
6 sections; and

7 a matrix describing said hardware accelerators and said designs
8 configured to support run time execution.

1 23. The system of claim 22, wherein said hardware accelerators includes

2 digital signal processors.

1 24. The system of claim 22, wherein said hardware accelerators includes

2 bins.

1 25. The system of claim 24, wherein said bins support multiple hardware

2 contexts.

1 26. The system of claim 25, wherein said bins support a first set of variants
2 configured to support said multiple hardware contexts.

1 27. The system of claim 26, wherein said first set of variants are produced
2 based upon resource usage.

1 28. The system of claim 27, wherein a second set of variants of said designs
2 are configured to support multiple hardware configurations of one of said
3 plurality of bins.

1 29. The system of claim 22, wherein said mapping is performed by a place
2 and route.

1 30. The system of claim 22, wherein said decomposing is performed
2 manually.

1 31. The system of claim 22, wherein said decomposing is performed by a
2 software profiler.

1 32. The system of claim 31, wherein said software profiler executes code
2 compiled from said program description, and monitors time consumed.

1 33. The system of claim 32, wherein said software profiler includes a set of
2 test data.

1 34. The system of claim 32, wherein said software profiler determines
2 functions that consume a significant portion of said time consumed.

1 35. The system of claim 31, wherein said software profiler is configured to
2 identify kernel sections by identifying regular structures.

1 36. The system of claim 31, wherein said software profiler is configured to
2 identify kernel sections by identifying sections with a limited number of inputs
3 and outputs.

1 37. The system of claim 31, wherein said software profiler is configured to
2 identify kernel sections by identifying sections with a limited number of
3 branches.

1 38. The system of claim 31, wherein said profiler identifies overhead
2 sections.

1 39. The system of claim 22, wherein said designs include microcode.

1 40. The system of claim 39, wherein said microcode includes context
2 dependent configurations.

1 41. The system of claim 22, wherein said matrix is sparsely-populated.

1 42. The system of claim 22, wherein said matrix is fully-populated.

1 43. A machine-readable medium having stored thereon instructions for
2 processing elements, which when executed by said processing elements
3 perform the following:
4 partitioning a processing element array into a plurality of hardware
5 accelerators;
6 decomposing a program description into a plurality of kernel sections;
7 mapping said kernel sections into a plurality of hardware dependent
8 designs; and
9 forming a matrix describing said hardware accelerators and said designs
10 configured to support run time execution.

1 44. A system configured to create run time executable code, comprising:
2 means for partitioning a processing element array into a plurality of
3 hardware accelerators;
4 means for decomposing a program description into a plurality of kernel
5 sections;
6 means for mapping said kernel sections into a plurality of hardware
7 dependent designs; and
8 means for forming a matrix describing said hardware accelerators and
9 said designs configured to support run time execution.

ABSTRACT

A system and method for creating run time executables in a configurable processing element array is disclosed. This system and method includes the

5 step of partitioning a processing element array into a number of defined sets of hardware accelerators, which in one embodiment are processing elements called “bins”. The system and method then involves decomposing a program description in object code form into a plurality of “kernel sections”, where the kernel sections are defined as those sections of object code which are

10 candidates for hardware acceleration. Next, mapping the identified kernel sections into a number of hardware dependent designs is performed. Finally, a matrix of the bins and the designs is formed for use by the run time system.

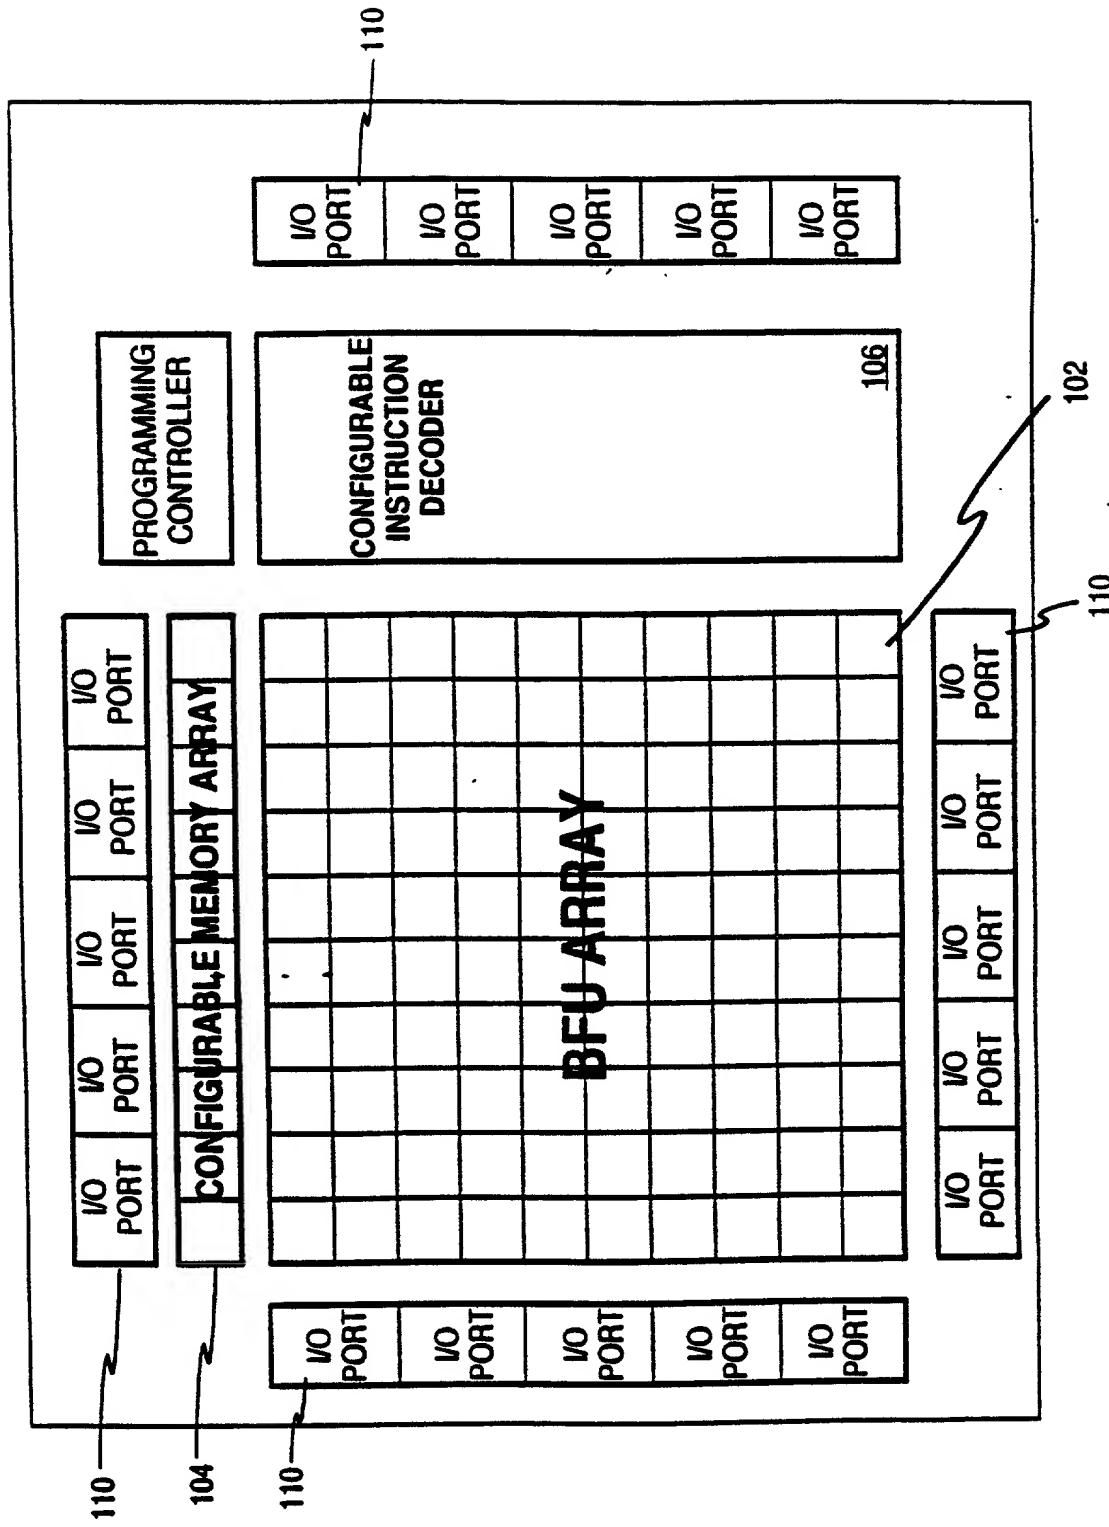


FIG. 1

MCPE

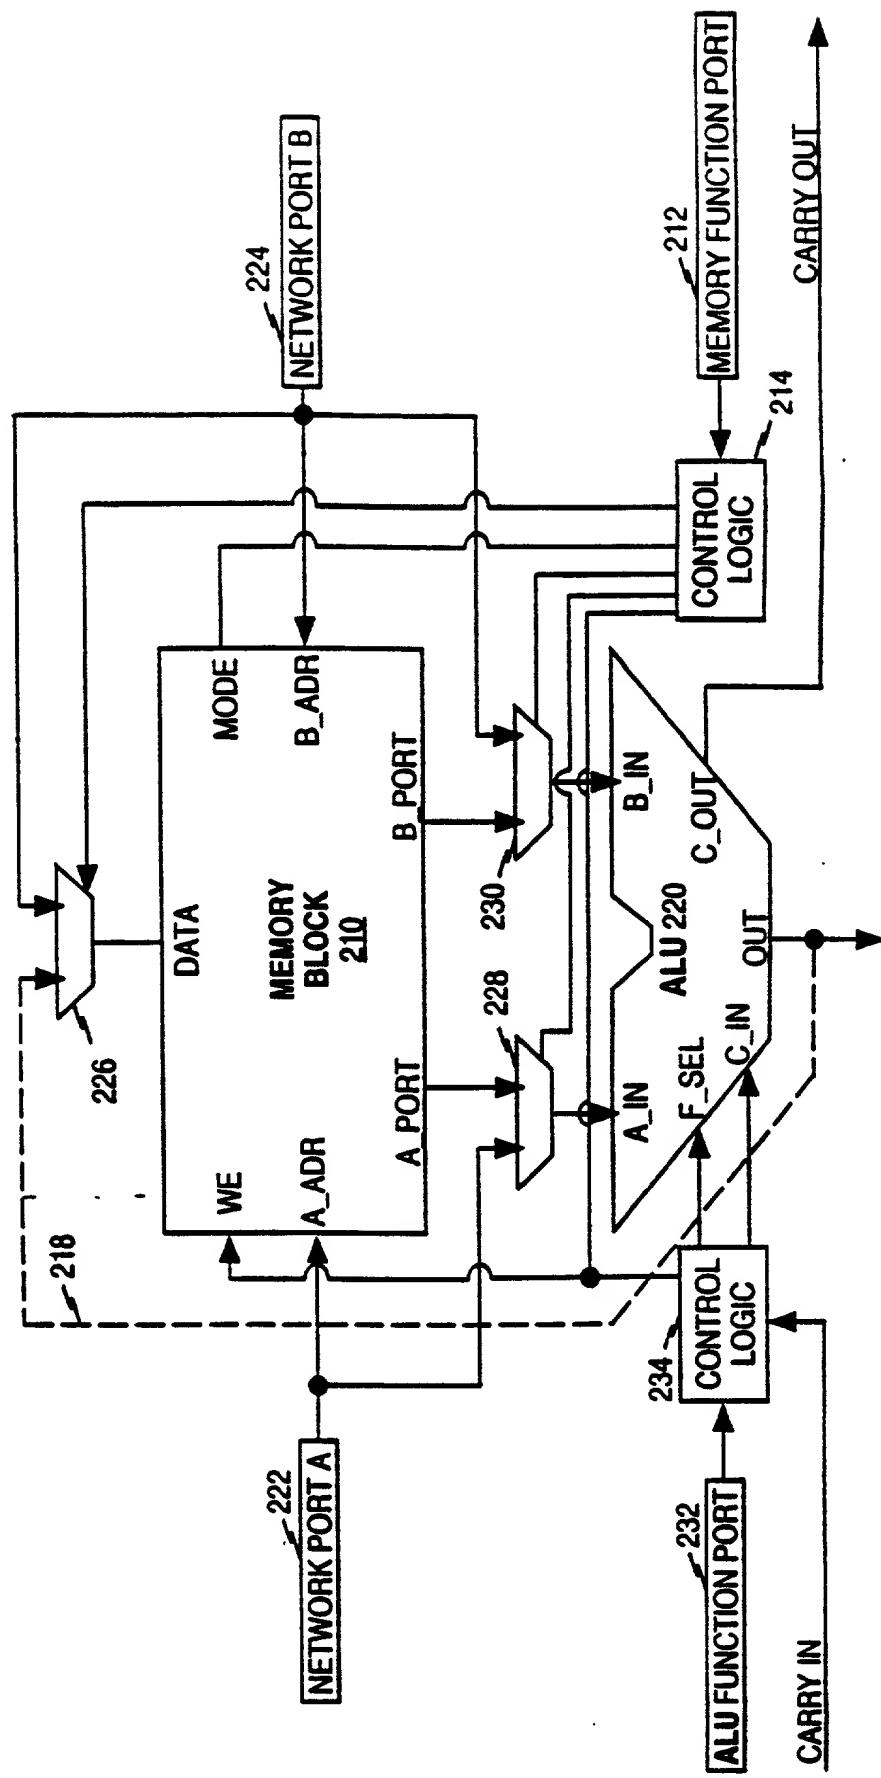


FIG. 2

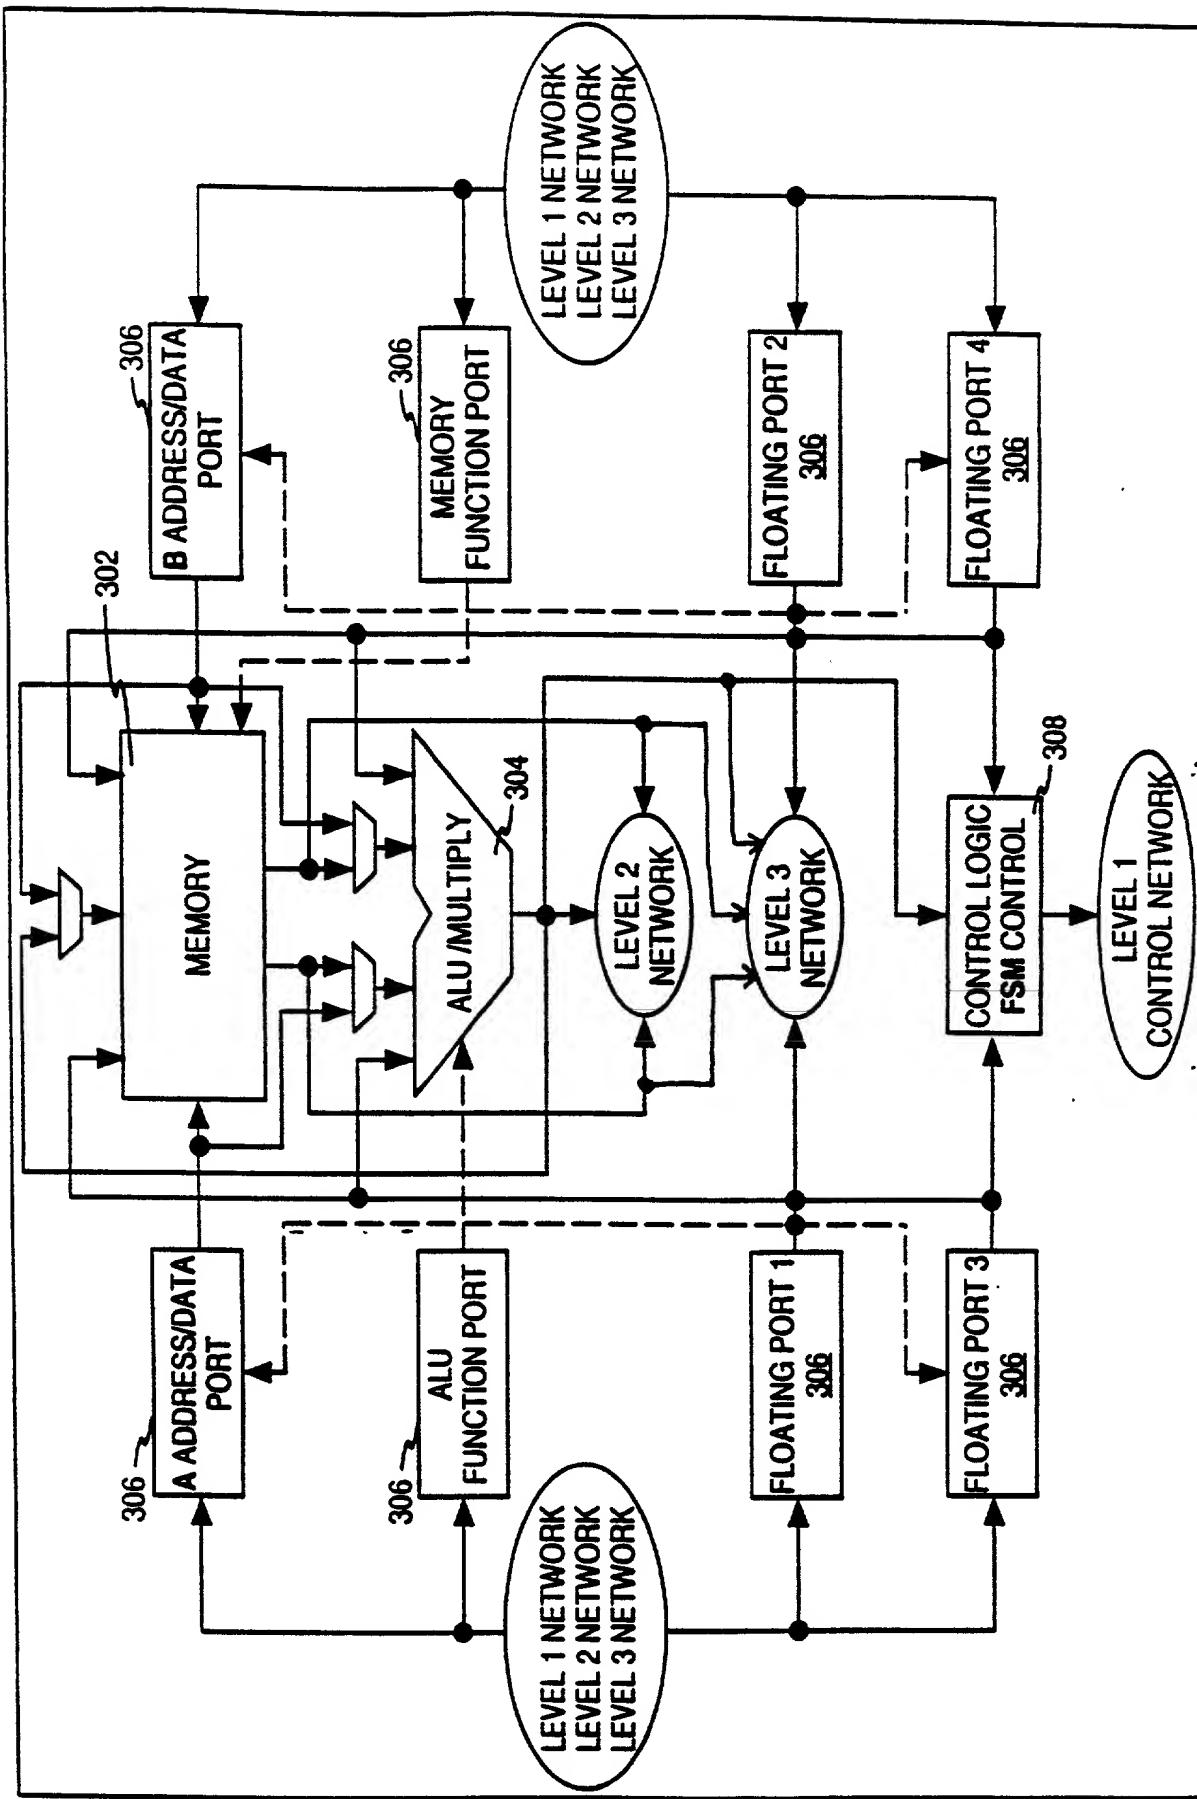


FIG. 3

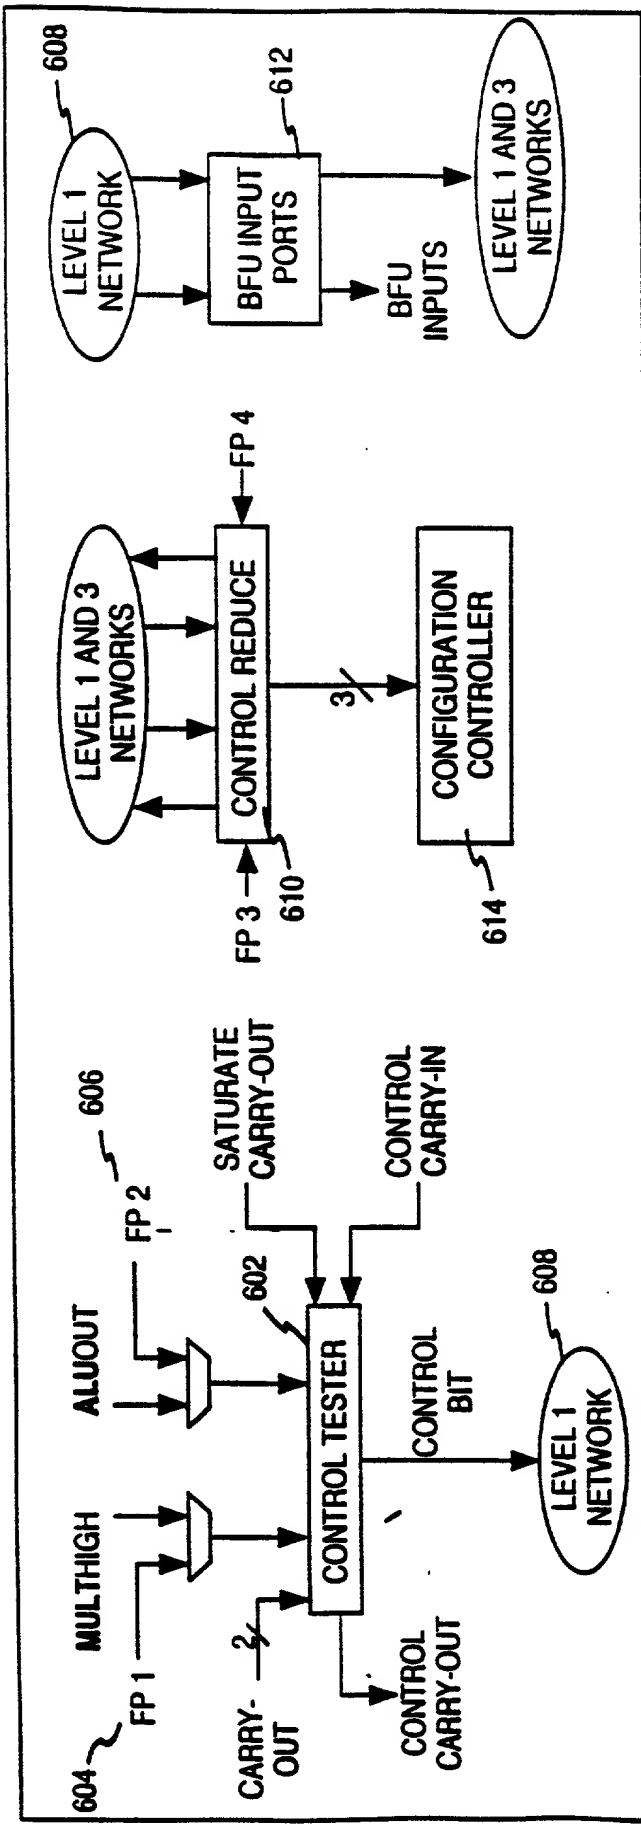


FIG. 4

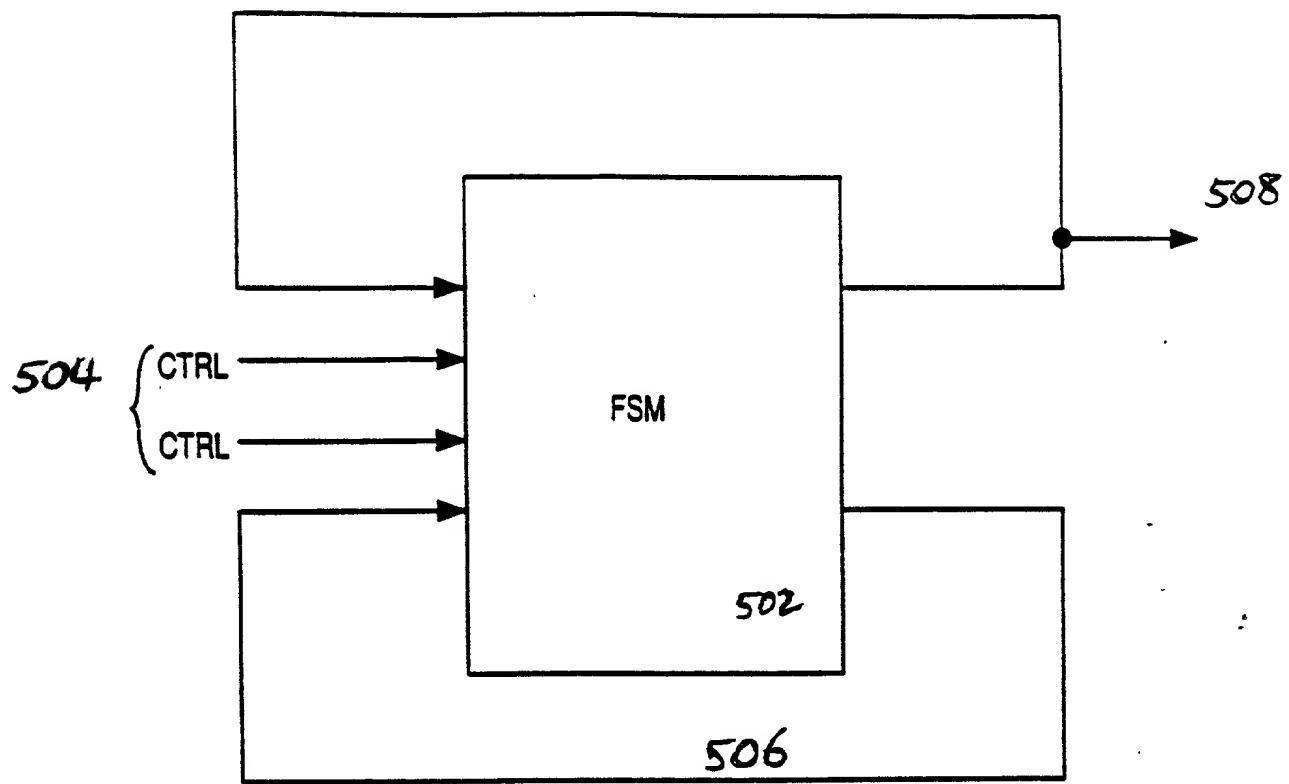


FIG. 5

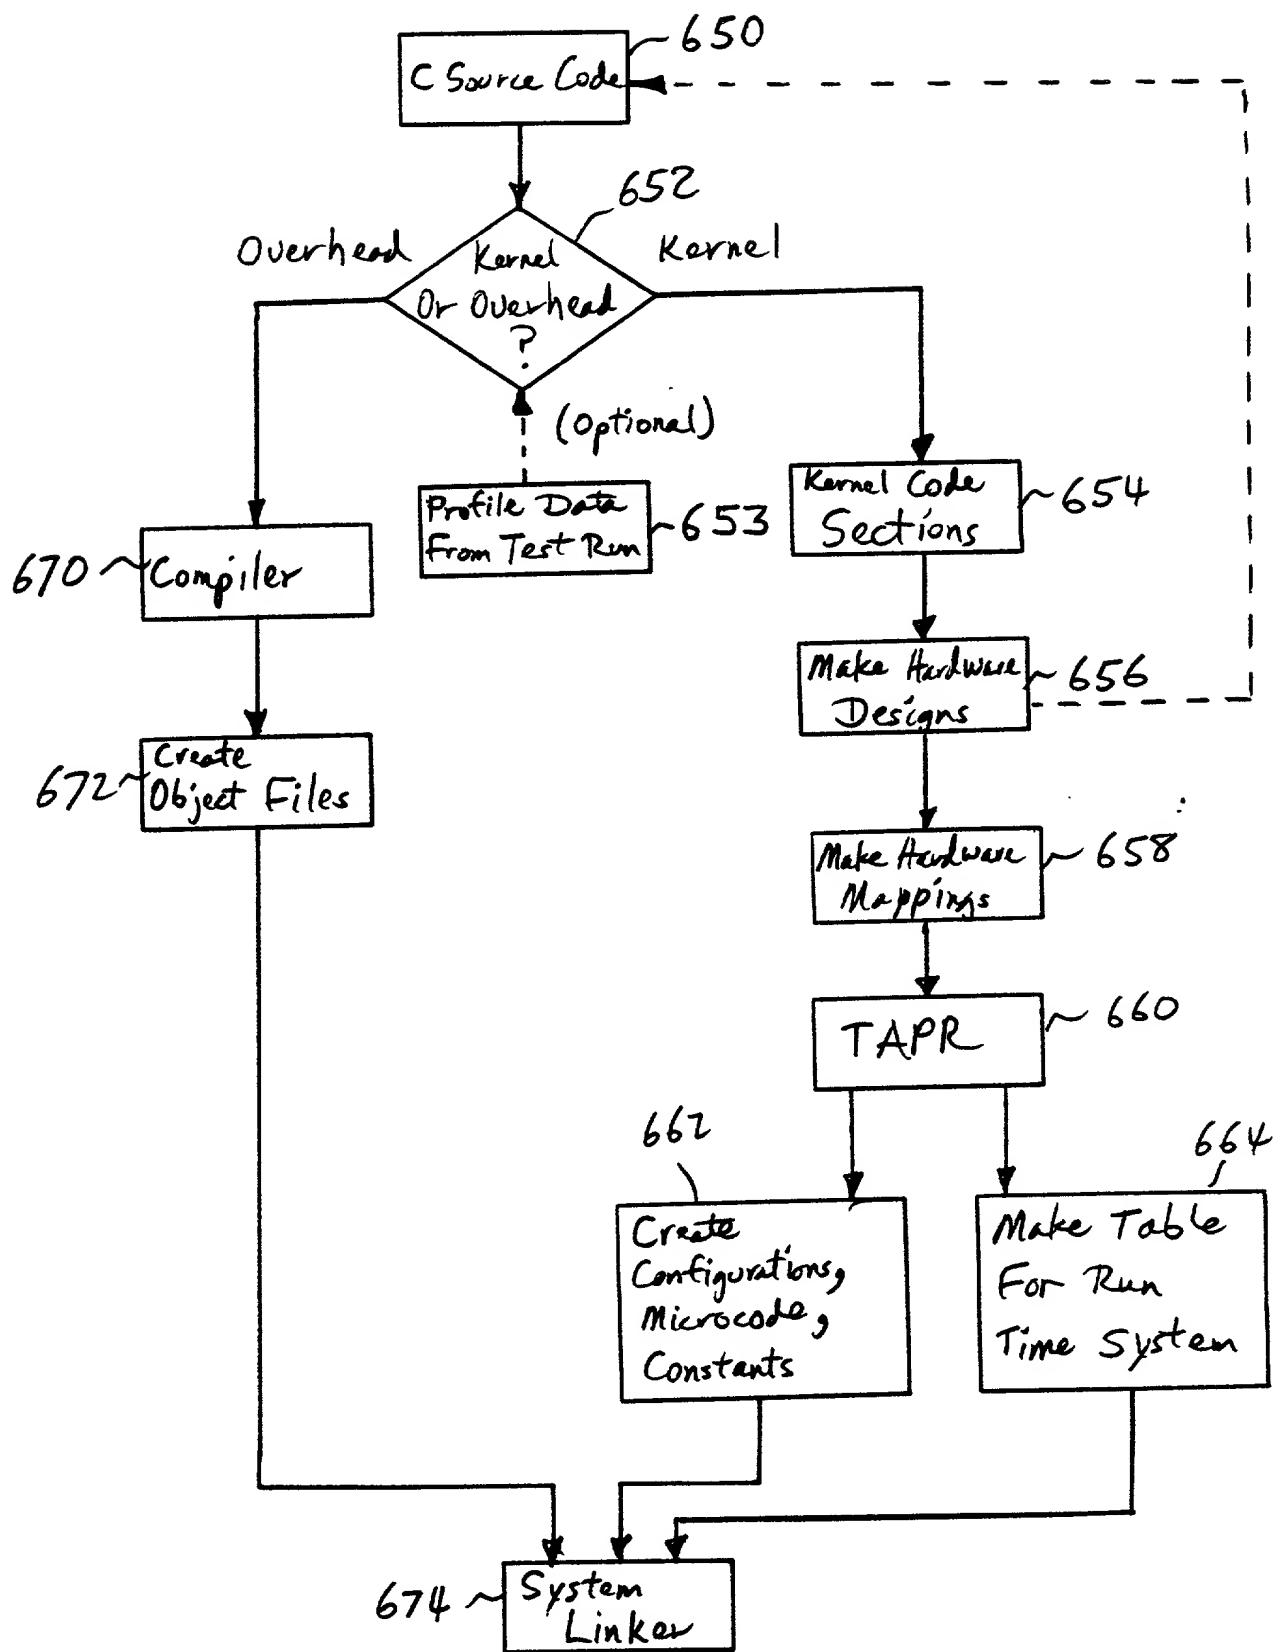


FIGURE 6

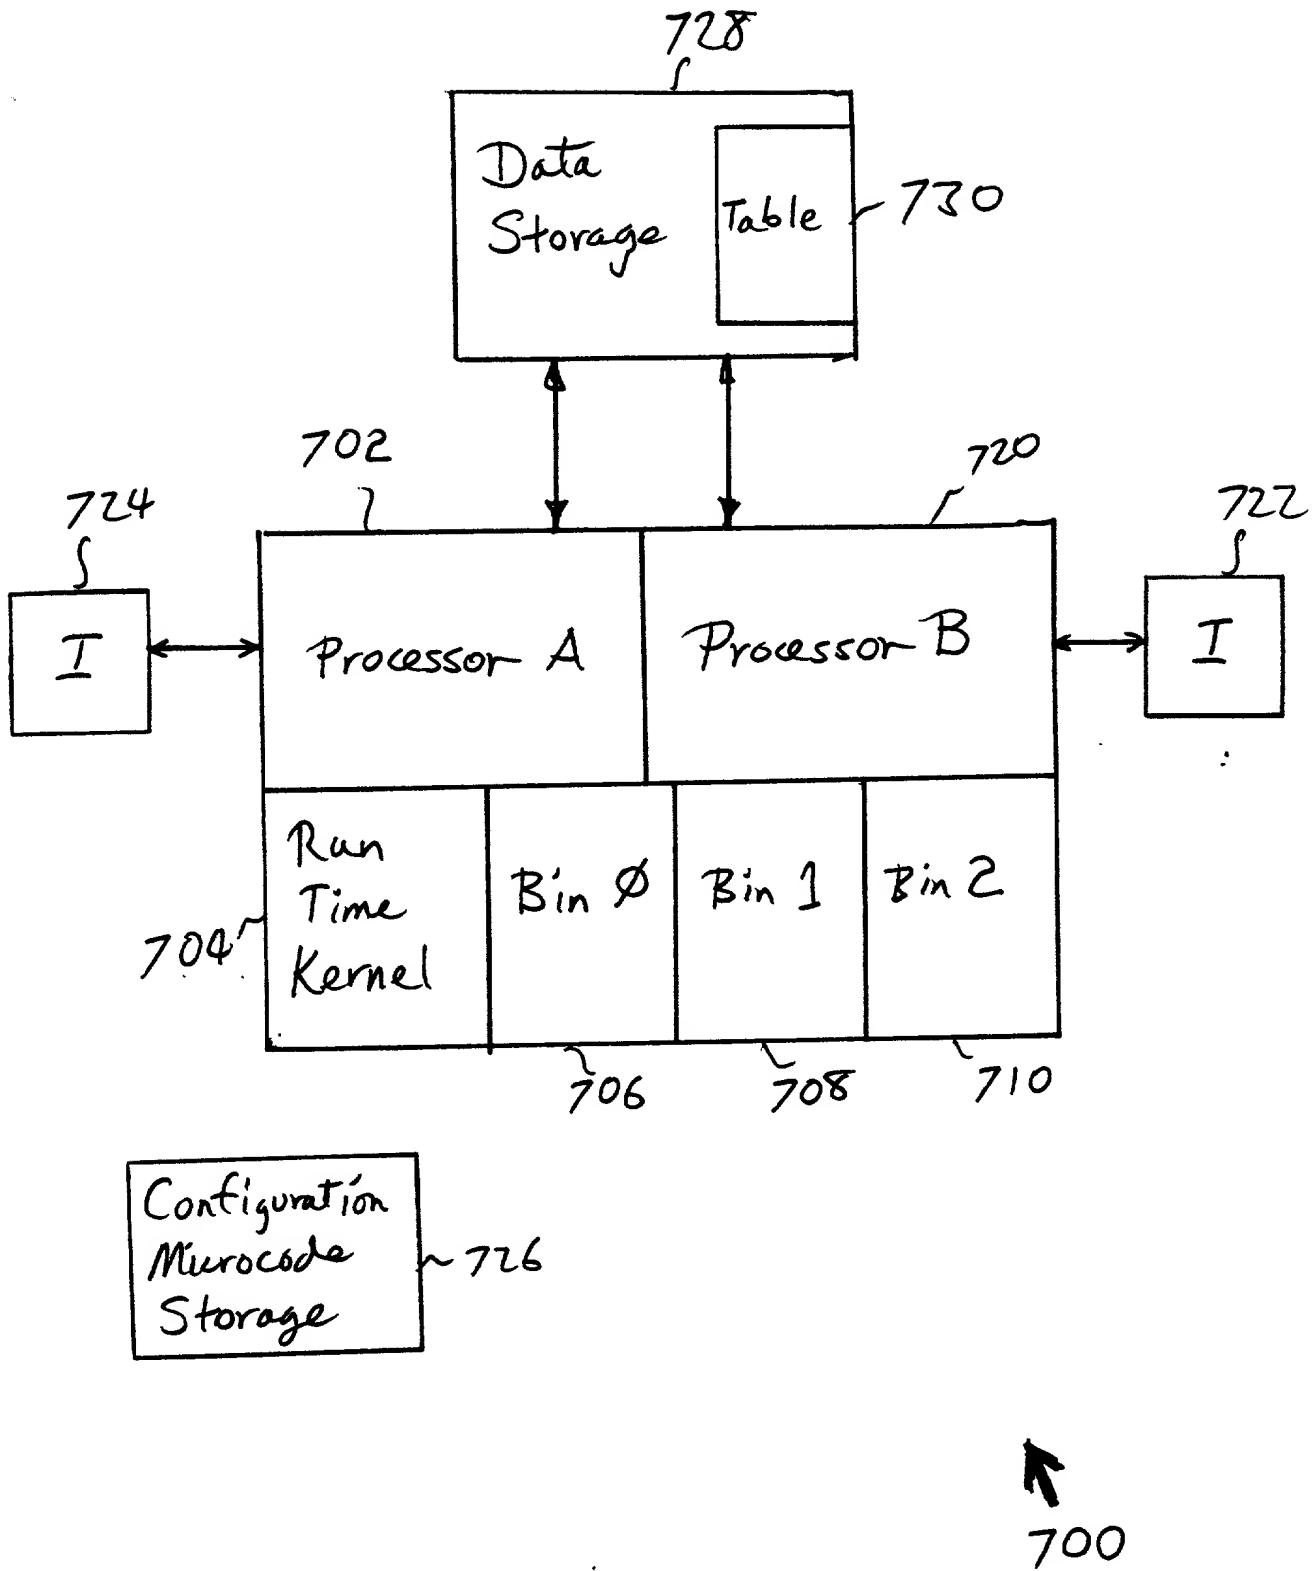
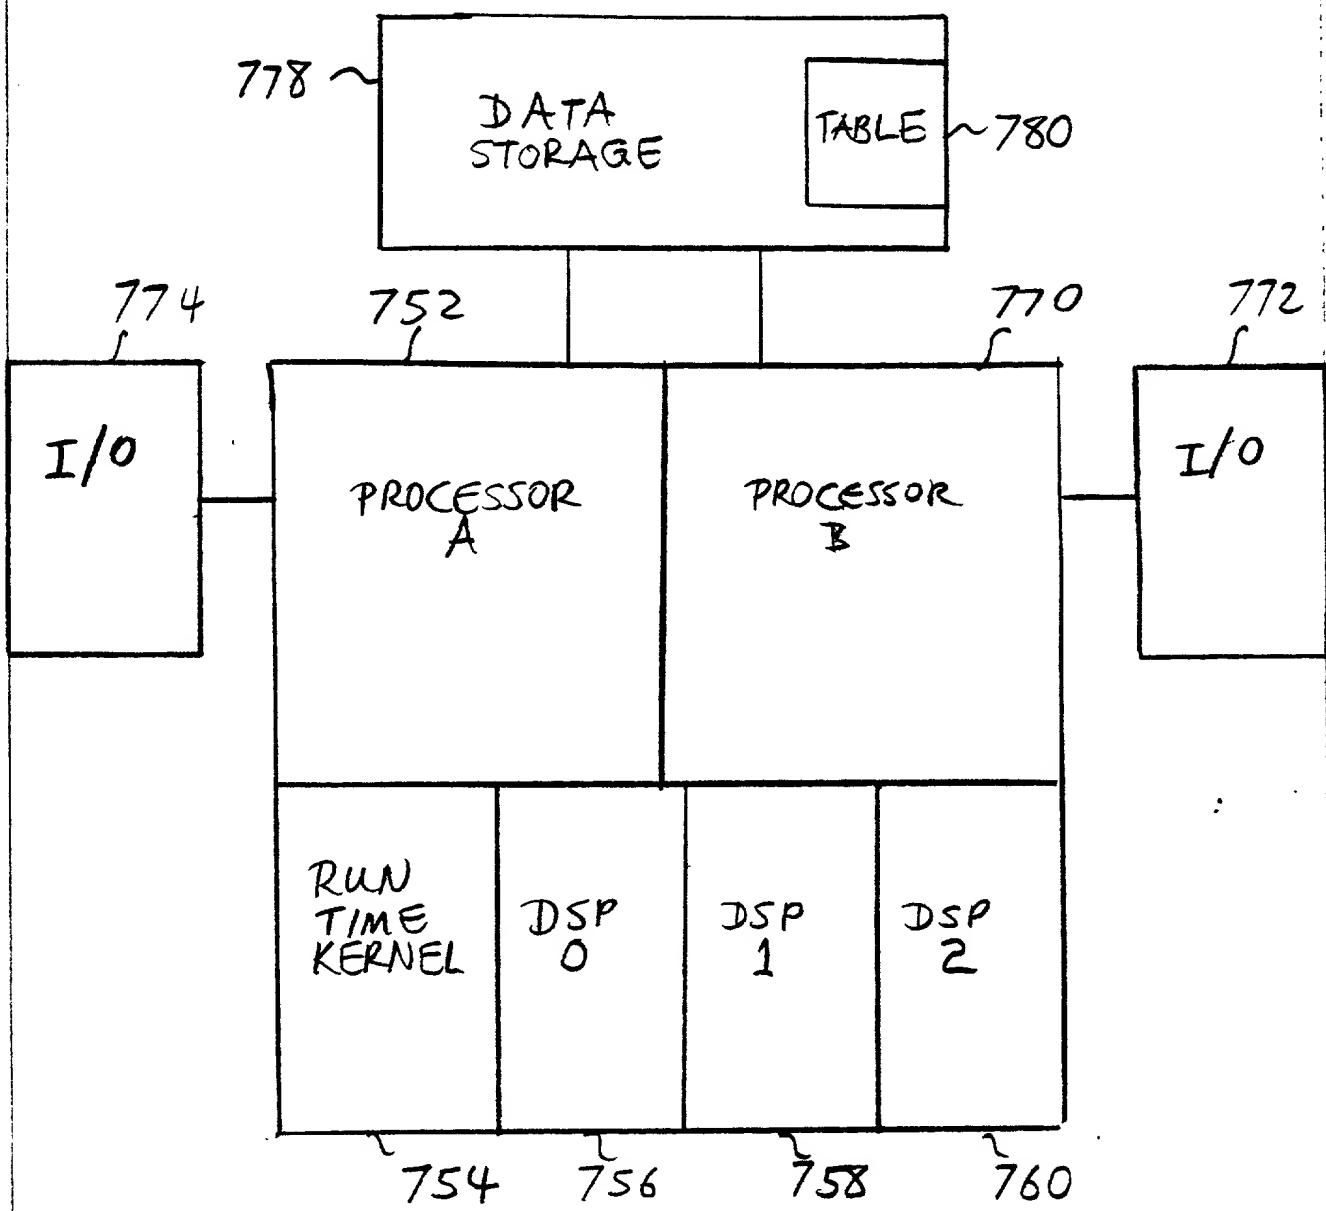


FIGURE 7A



750

FIG 7B

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



00000000000000000000000000000000

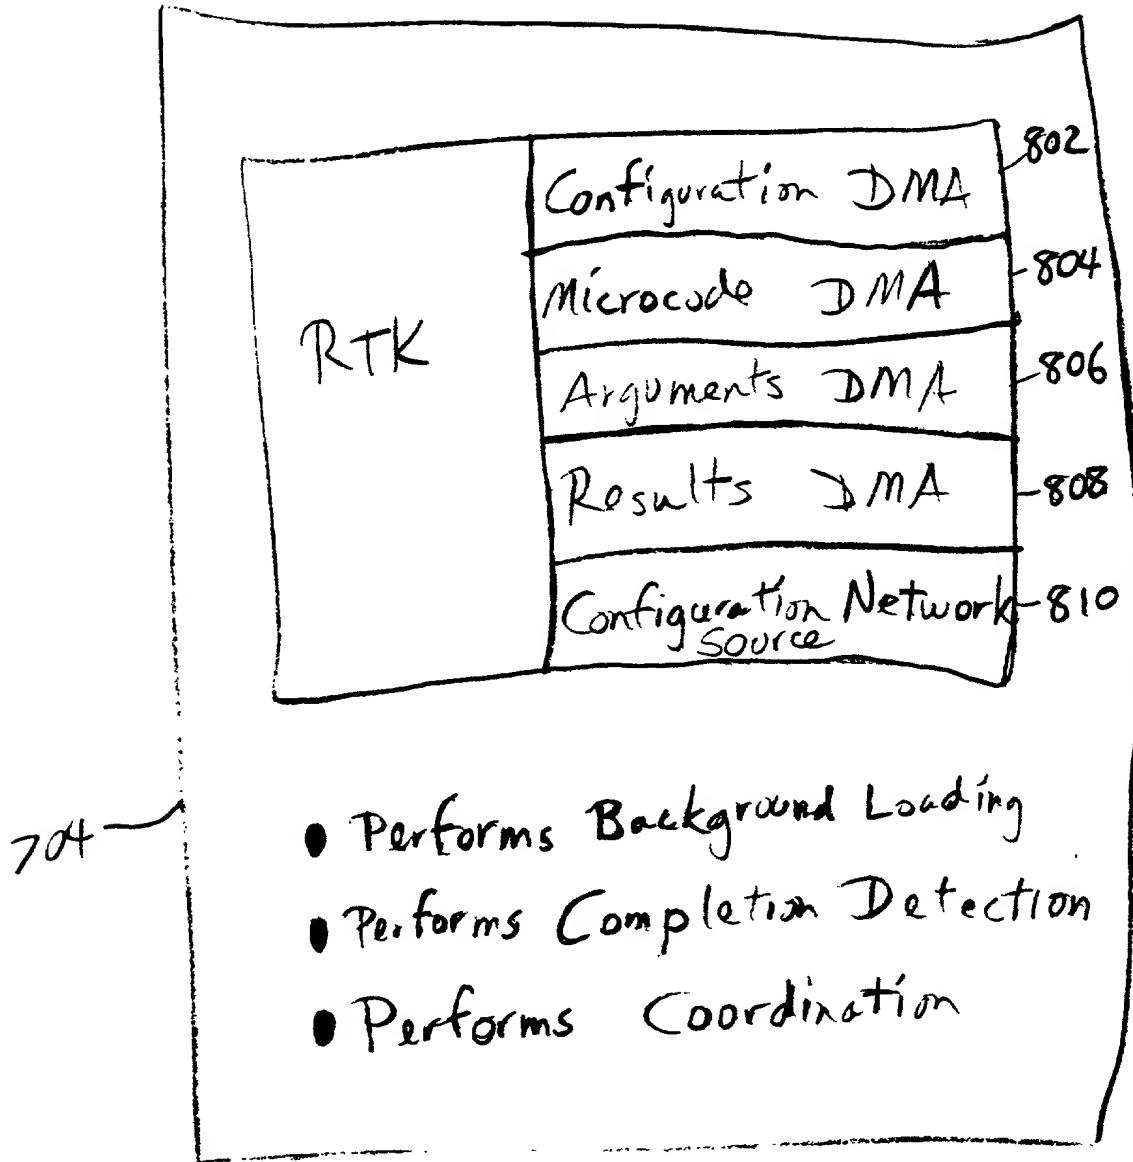


FIG.8

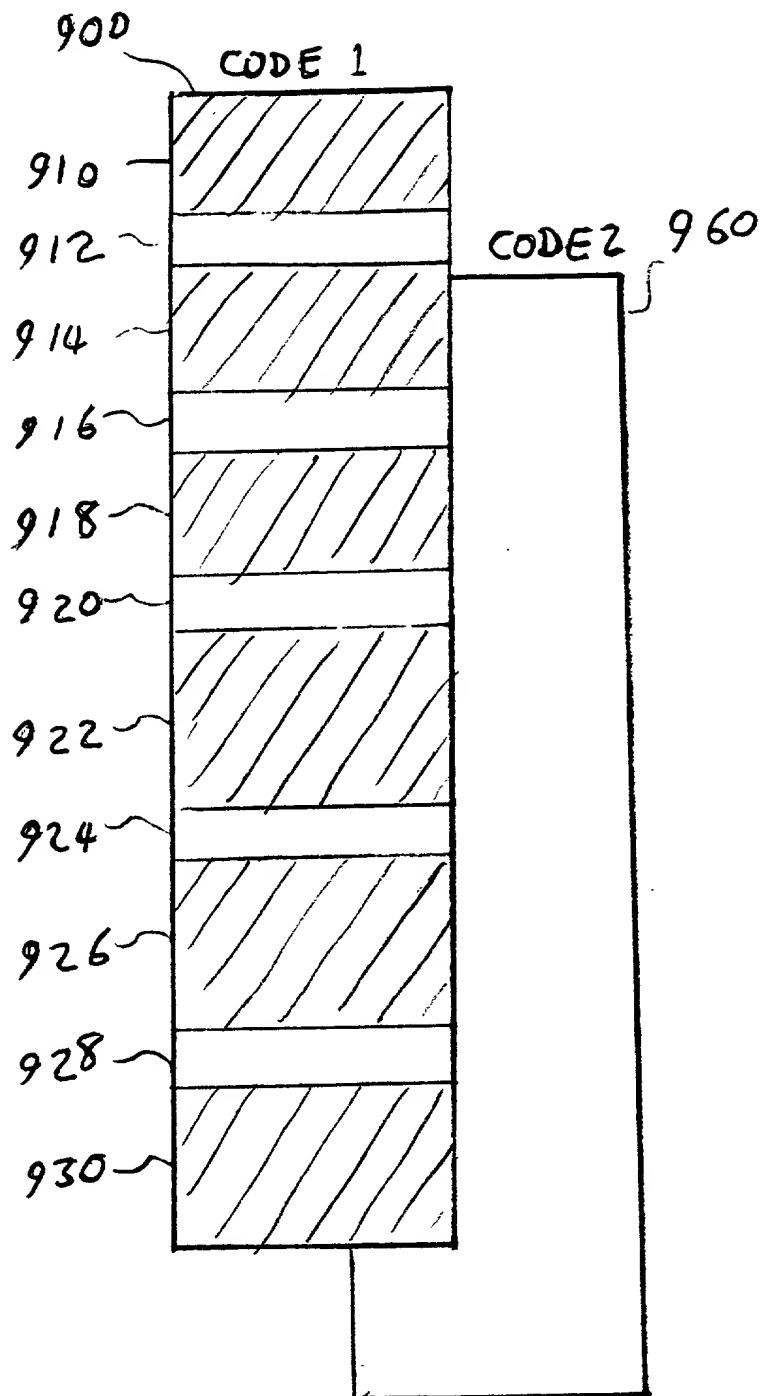


FIG. 9A

000 000 000 000 000 000

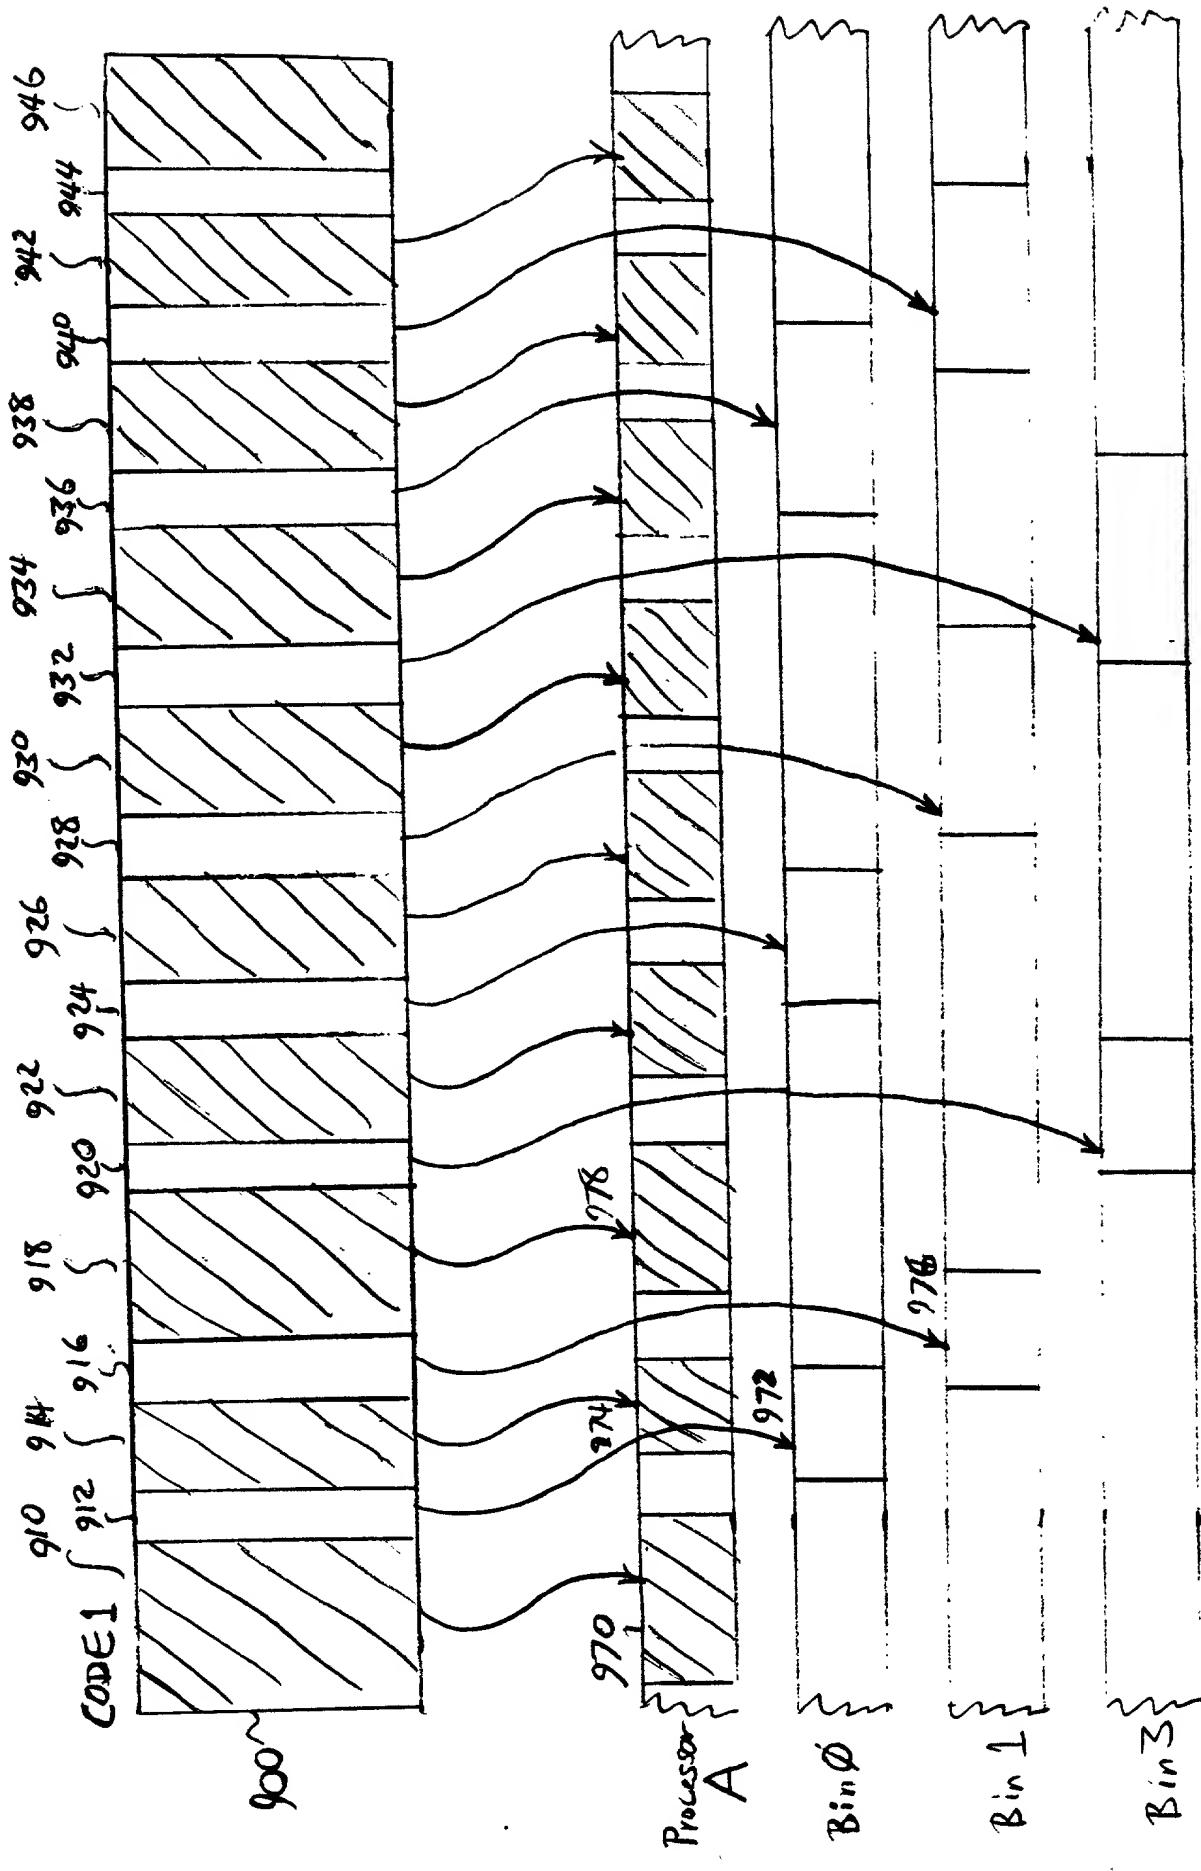


FIG 9B

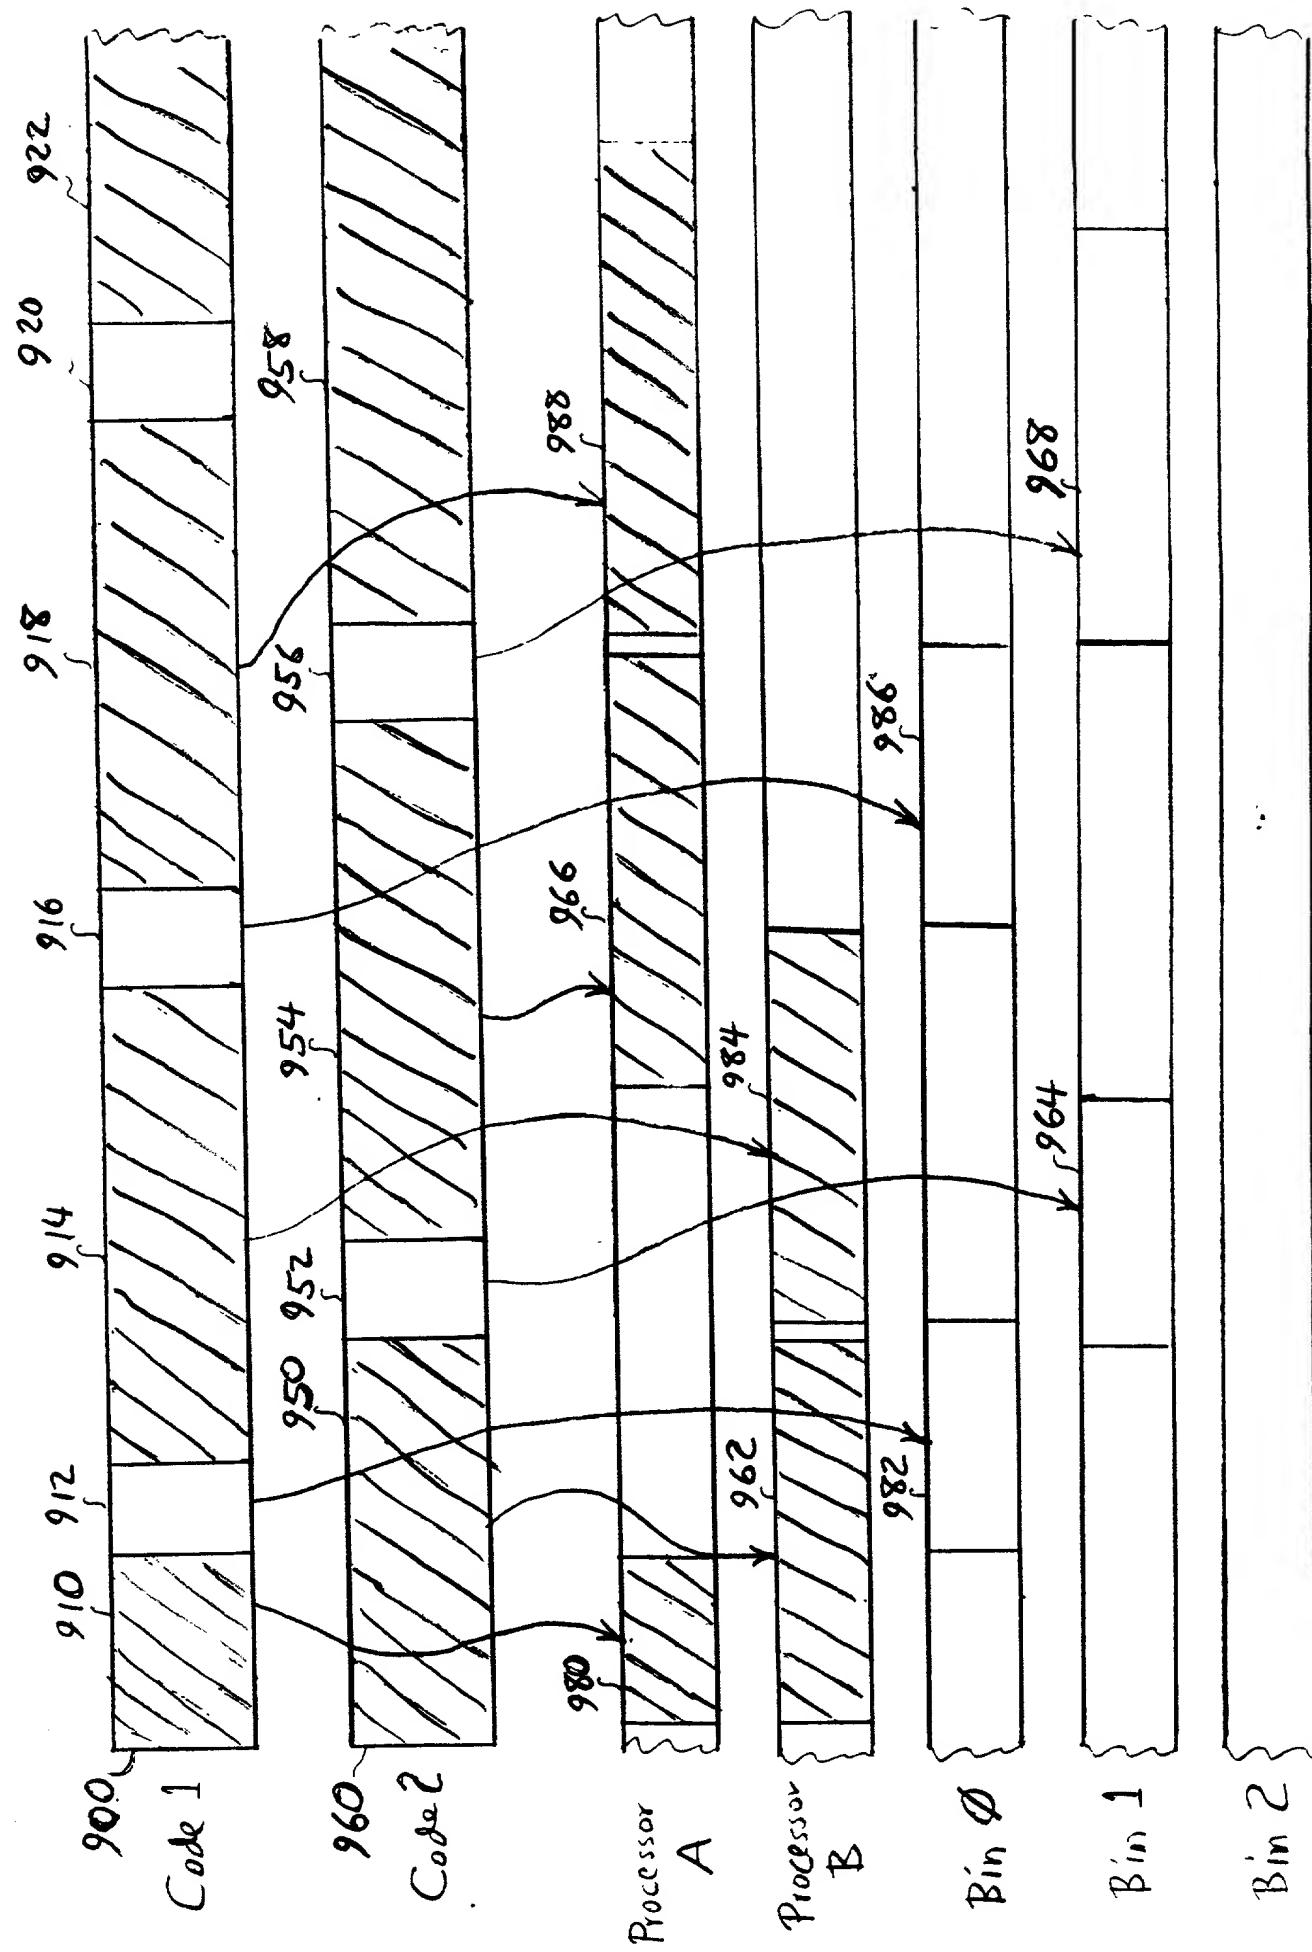


FIG. 9C

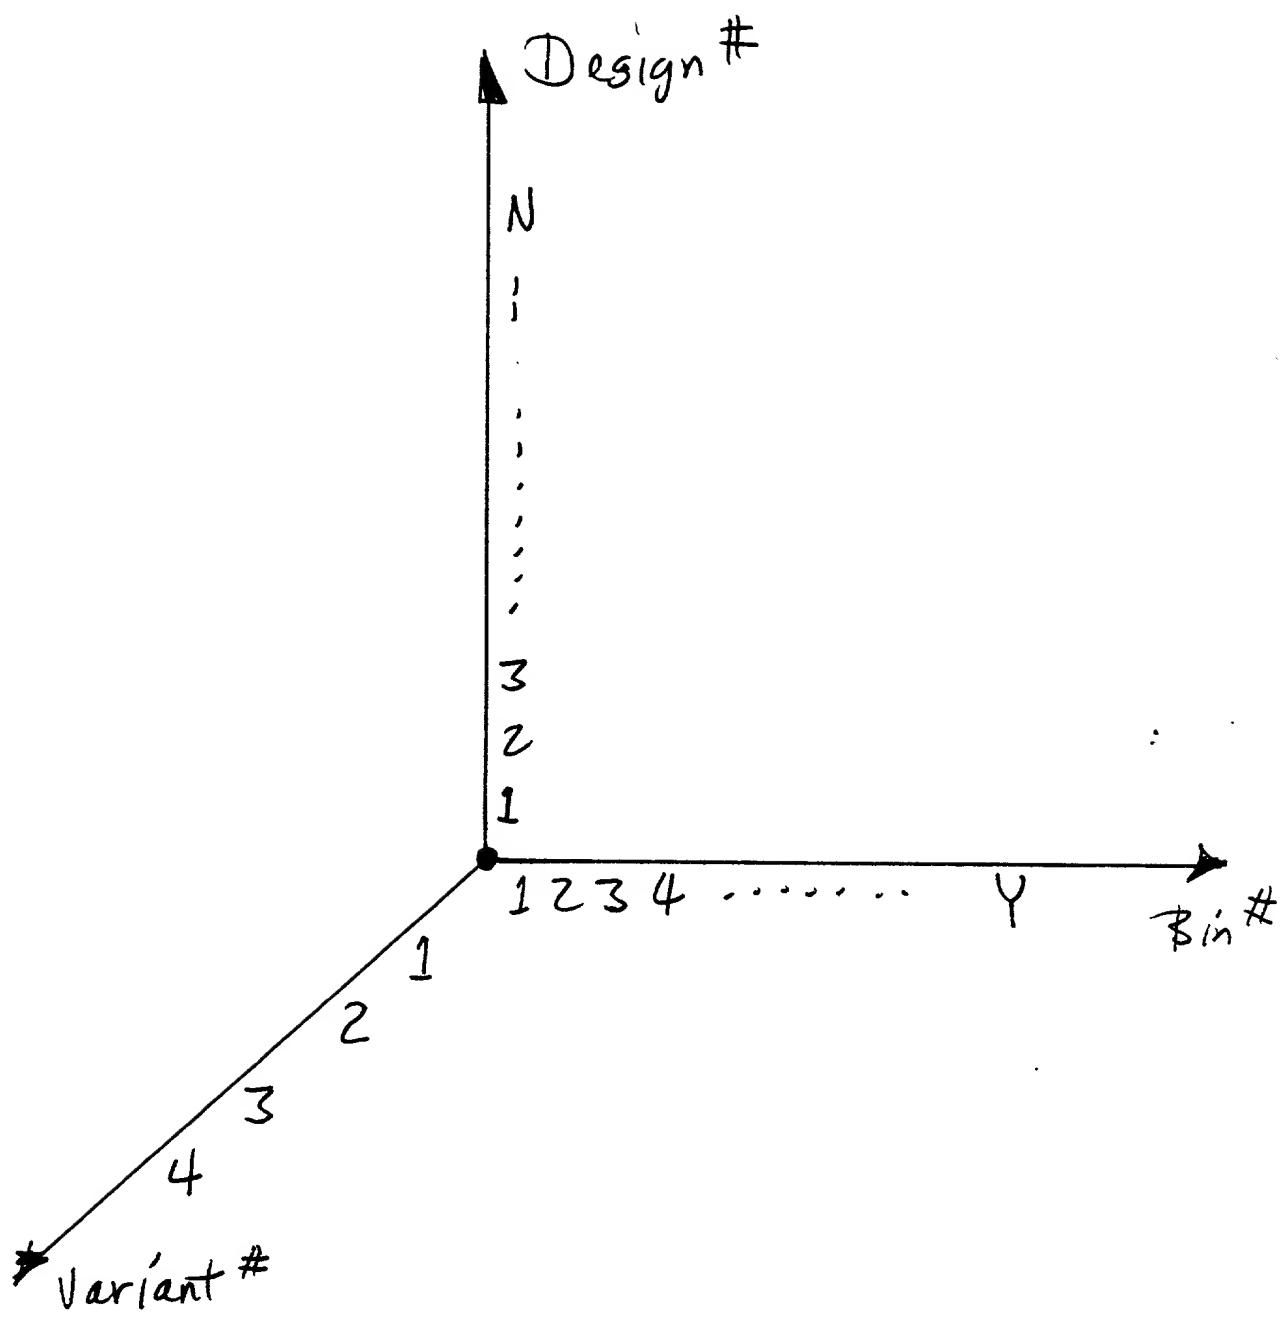


FIGURE 10

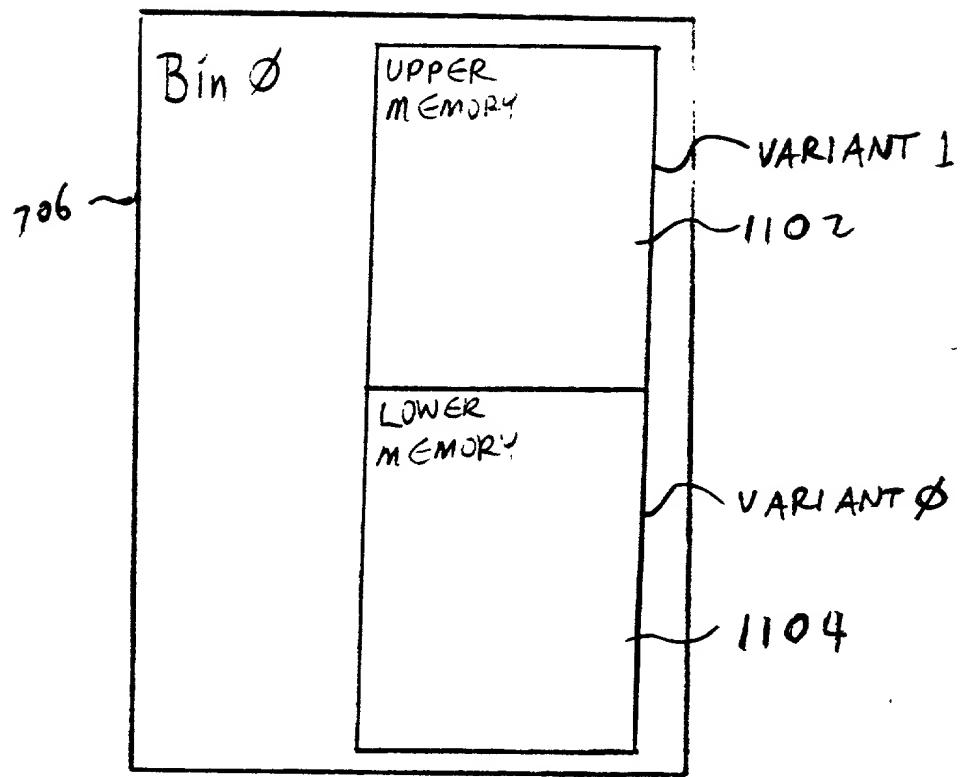


FIG. 11

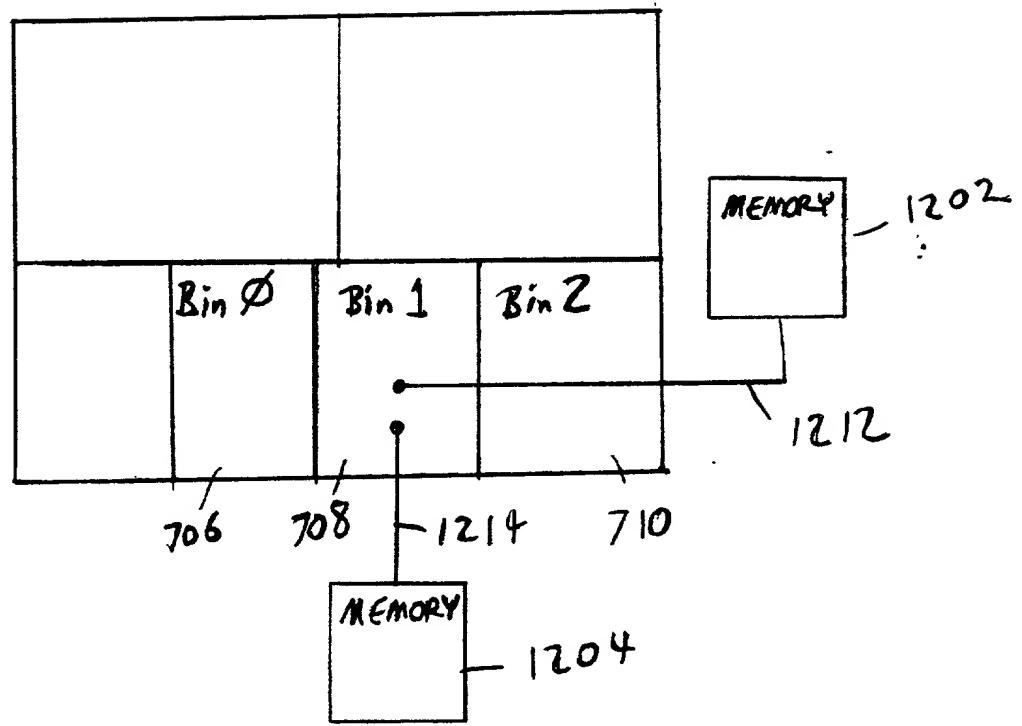


FIG. 12

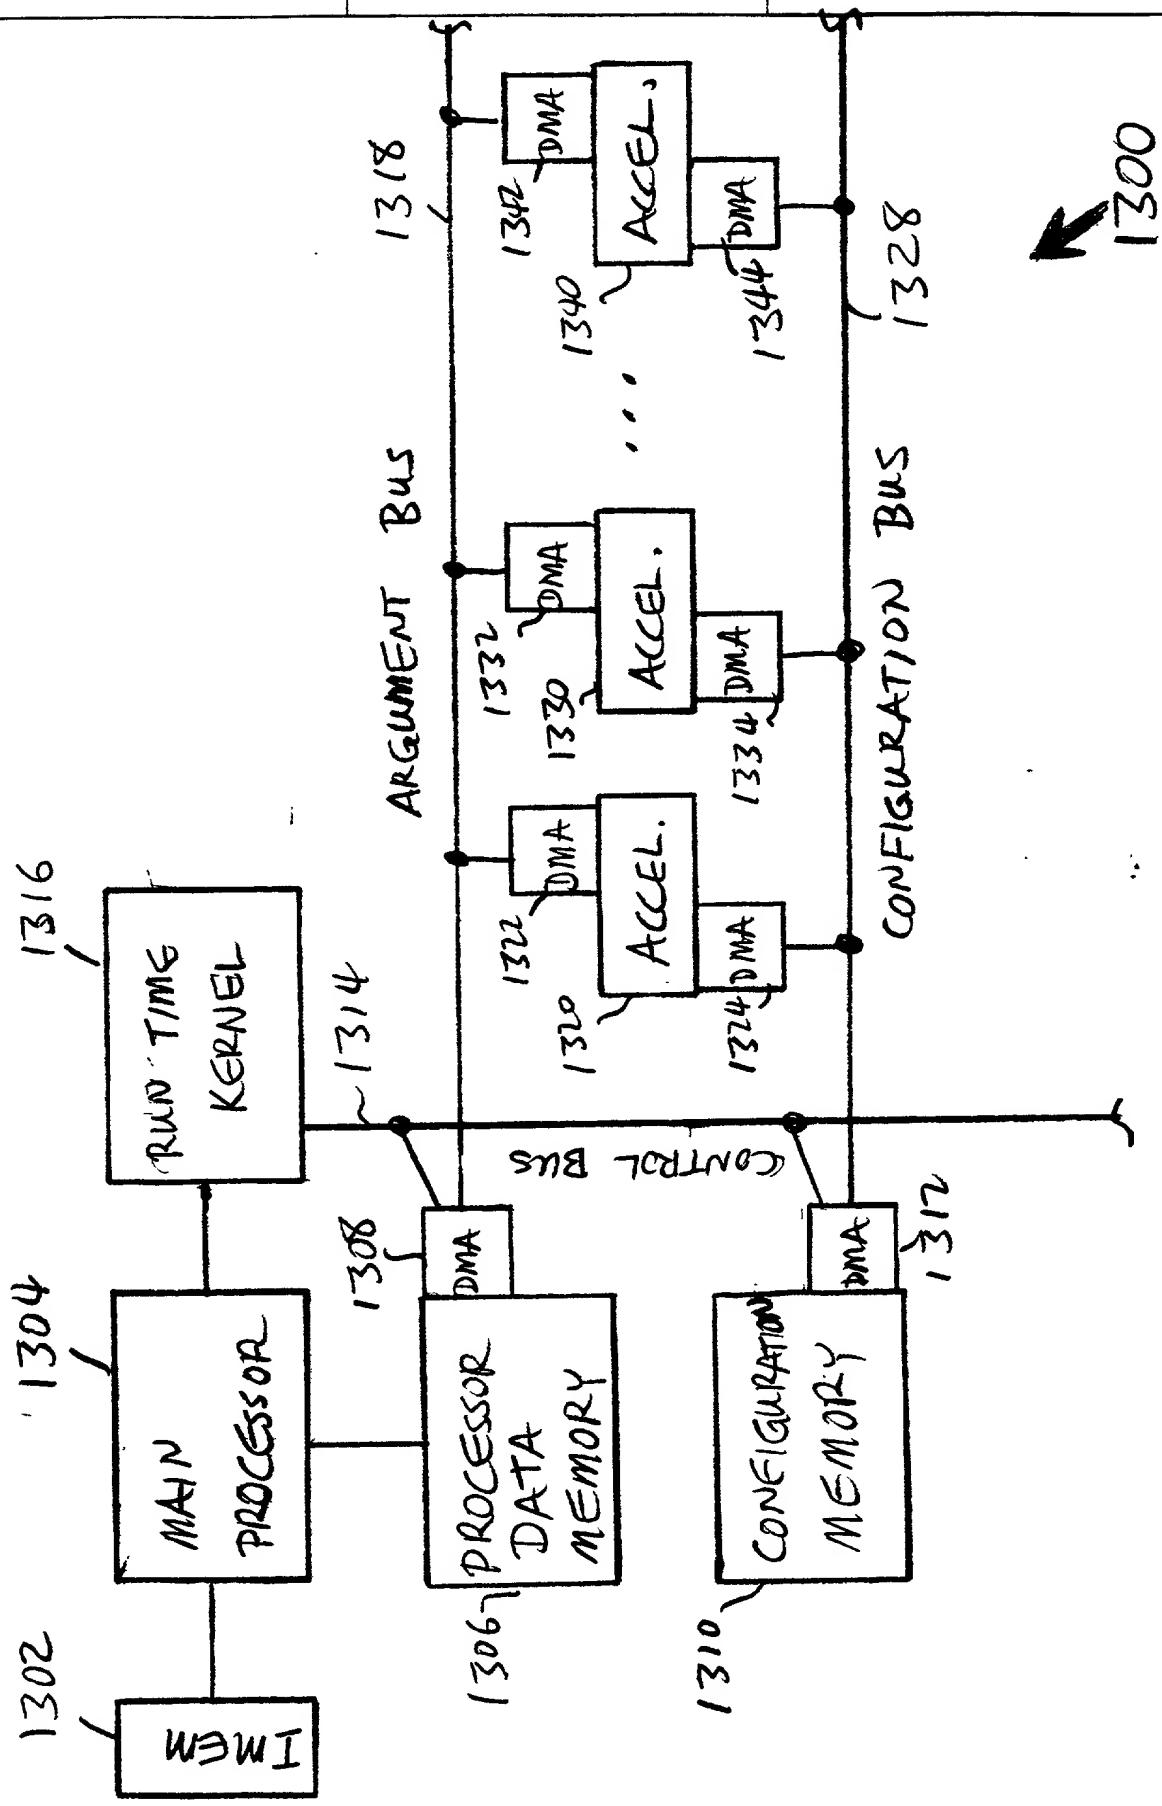


FIG 13

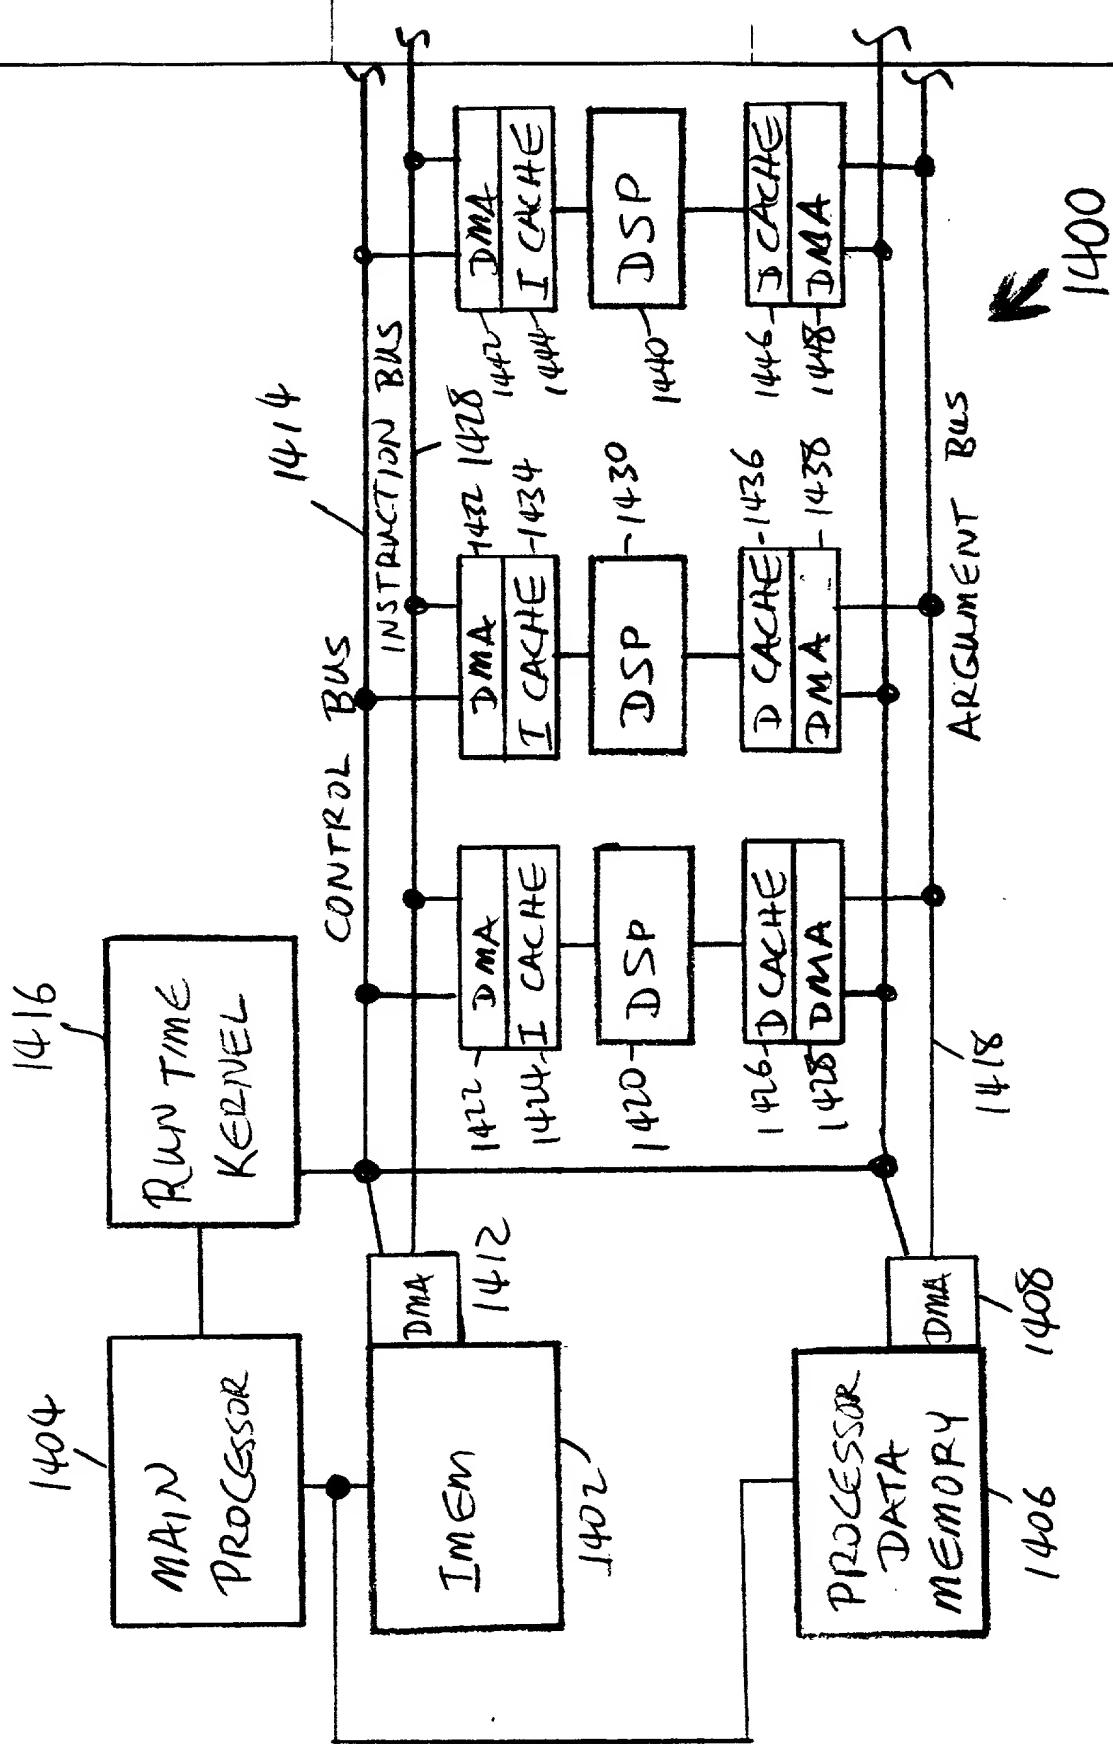


FIG 14

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

SYSTEM AND METHOD FOR PREPARING SOFTWARE FOR EXECUTION IN A DYNAMICALLY CONFIGURABLE HARDWARE ENVIRONMENT

the specification of which

X is attached hereto.
 _____ was filed on _____ as
 United States Application Number _____
 or PCT International Application Number _____
 and was amended on _____
 (if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>	<u>Priority Claimed</u>			
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No

I hereby claim the benefit under title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below:

(Application Number)	Filing Date
----------------------	-------------

(Application Number)	Filing Date
----------------------	-------------

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Number)	Filing Date	(Status -- patented, pending, abandoned)
----------------------	-------------	---

(Application Number)	Filing Date	(Status -- patented, pending, abandoned)
----------------------	-------------	---

I hereby appoint the persons listed on Appendix A hereto (which is incorporated by reference and a part of this document) as my respective patent attorneys and patent agents, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to Dennis A. Nicholls, BLAKELY, SOKOLOFF, TAYLOR &
(Name of Attorney or Agent)
ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct
telephone calls to Dennis A. Nicholls, (408) 720-8300.
(Name of Attorney or Agent)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Christopher Songer

Inventor's Signature C. M. Songer Date 24 Aug 2000

Residence Sunnyvale, California Citizenship USA

(City, State)

(Country)

Post Office Address 693 Windsor Terrace

Sunnyvale, California 94087

Full Name of Second/Joint Inventor Ian S. Eslick

Inventor's Signature C. M. Eslick Date 8/25/2000

Residence Mountain View, California Citizenship USA

(City, State)

(Country)

Post Office Address 652 California Street #A

Mountain View, California 94041

Full Name of Third/Joint Inventor Robert S. French

Inventor's Signature R. S. French Date 8/27/00

Residence Sunnyvale, California Citizenship USA

(City, State)

(Country)

Post Office Address 1712 Kimberly Drive

Sunnyvale, California 94087

APPENDIX A

William E. Alford, Reg. No. 37,764; Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; Gregory D. Caldwell, Reg. No. 39,926; Ronald C. Card, Reg. No. 44,587; Andrew C. Chen, Reg. No. 43,544; Thomas M. Coester, Reg. No. 39,637; Alin Corie, Reg. No. P46,244; Dennis M. deGuzman, Reg. No. 41,702; Stephen M. De Clerk, under 37 C.F.R. § 10.9(b); Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Sanjeet Dutta, Reg. No. P46,145; Matthew C. Fagan, Reg. No. 37,542; Tarek N. Fahmi, Reg. No. 41,402; Paramita Ghosh, Reg. No. 42,806; James Y. Go, Reg. No. 40,621; James A. Henry, Reg. No. 41,064; Willmore F. Holbrow III, Reg. No. P41,845; Sheryl Sue Holloway, Reg. No. 37,850; George W Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; William W. Kidd, Reg. No. 31,772; Sang Hui Kim, Reg. No. 40,450; Eric T. King, Reg. No. 44,188; Erica W. Kuo, Reg. No. 42,775; Kurt P. Leyendecker, Reg. No. 42,799; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Darren J. Milliken, Reg. 42,004; Lisa A. Norris, Reg. No. 44,976; Chun M. Ng, Reg. No. 36,878; Thien T. Nguyen, Reg. No. 43,835; Thinh V. Nguyen, Reg. No. 42,034; Dennis A. Nicholls, Reg. No. 42,036; Daniel E. Ovanezian, Reg. No. 41,236; Marina Portnova, Reg. No. P45,750; Babak Redjaian, Reg. No. 42,096; William F. Ryann, Reg. 44,313; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Jeffrey Sam Smith, Reg. No. 39,377; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; John F. Travis, Reg. No. 43,203; George G. C. Tseng, Reg. No. 41,355; Joseph A. Twarowski, Reg. No. 42,191; Lester J. Vincent, Reg. No. 31,460; Glenn E. Von Tersch, Reg. No. 41,364; John Patrick Ward, Reg. No. 40,216; Mark L. Watson, Reg. No. P46,322; Thomas C. Webster, Reg. No. P46,154; Charles T. J. Weigell, Reg. No. 43,398; Kirk D. Williams, Reg. No. 42,229; James M. Wu, Reg. No. 45,241; Steven D. Yates, Reg. No. 42,242; and Norman Zafman, Reg. No. 26,250; my patent attorneys, and Justin M. Dillon, Reg. No. 42,486; my patent agent, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and James R. Thein, Reg. No. 31,710, my patent attorney.

APPENDIX B

Title 37, Code of Federal Regulations, Section 1.56 Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclosure information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclosure all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) Prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.